EUROPEAN CENTRAL BANK
EUROSYSTEM

Working Paper Series

Giovanni Bonfanti

Fiscal integration as a Pareto improvement

ECB – Lamfalussy Fellowship Programme

No 3297 / October 2026

Disclaimer: This paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

ECB Lamfalussy Fellowship Programme

This paper has been produced under the ECB Lamfalussy Fellowship programme. This programme was launched in 2003 in the context of the ECB-CFS Research Network on “Capital Markets and Financial Integration in Europe”. It aims at stimulating high-quality research on the structure, integration and performance of the European financial system.

The Fellowship programme is named after Baron Alexandre Lamfalussy, the first President of the European Monetary Institute. Mr Lamfalussy is one of the leading central bankers of his time and one of the main supporters of a single capital market within the European Union.

Each year the programme sponsors five young scholars conducting a research project in the priority areas of the Network. The Lamfalussy Fellows and their projects are chosen by a selection committee composed of Eurosystem experts and academic scholars. Further information about the Network can be found at http://www.eufinancial-system.org and about the Fellowship programme under the menu point “fellowships”.

Abstract

Fiscal integration in a currency union is generally understood as a transfer from fiscally strong to fiscally weak countries. I show that this changes when common debt is paired with a non-defaultable common tax, even if national debt remains defaultable. The risky country gains from intertemporal trade: it receives a grant financed by a perpetual obligation it discounts more heavily than lenders do. The cost to the safe member is second-order and the tax removes it by insuring imperfectly correlated cycles rather than transferring resources in expectation. The calibration delivers a narrow Pareto-improving region: the tax that eliminates the risky country's perverse incentive to default is negligible, while the tax that compensates the safe country is an order of magnitude larger. Joint spending on a public good expands the region and reverses that ordering: the safe country gains even before the common budget balances.

JEL: F34, F36, F45, H63, H87.

Keywords: sovereign default, fiscal union, common debt, non-defaultable taxation, Pareto improvement.

Non-technical summary

The European Union has issued roughly €1 trillion of common debt since the start of the pandemic, and the 2024 Draghi report called for a permanent common borrowing capacity. The obstacle is political: fiscally strong member states read any common program as a transfer to fiscally weak ones, and no national government can support a program that leaves its own voters worse off.

This paper shows that the transfer reading need not hold. I study a currency area with a risky country that may default on its national debt and a safe country that does not, and add a common institution that issues bonds at the safe interest rate, distributes the proceeds to both members as an upfront grant, and services them with a proportional income tax levied on both. The tax is the essential feature: it is collected even in states in which the risky member defaults on its own bonds. National debt stays defaultable, and the institution neither guarantees nor retires it.

Both members can gain, and the gains are not driven by an expected transfer from one to the other. The risky member gains because it is impatient. It receives resources today and values the grant more than the payments that fund it. The gain does not come from cheaper national credit: a larger program leaves its own borrowing terms slightly worse rather than better. The safe member is compensated by insurance. Because the institution taxes both members and rebates revenue equally, resources flow each period toward whichever member draws the weaker income. They cancel in expectation, so neither country pays the other on average, but they arrive when they are most valuable, including when the risky member is in default.

Two tax rates matter and they are an order of magnitude apart. The first is the rate above which the common budget never runs a deficit. Below it, a defaulting member escapes its share of the shortfall, which strengthens the incentive to default where it is already strongest. For a program worth 10% of annual output that rate is about 0.09% of income. Clearing it, however, only removes the deficit channel; it does not compensate the safe member for accepting the program. Paying the safe member enough to accept it takes about 1.2%, thirteen times as much. Once the common tax exists, enforcing it is the easy part; compensating the creditor country is the expensive one.

The region of mutual gains is real but narrow, covering programs up to roughly 12.5% of annual output at modest tax rates. Three considerations enlarge it. Spending part of the common budget on a jointly consumed public good pays both members at once instead of moving resources between them, so the safe member gains before the budget balances. Uncertainty about which country will turn out to be risky works similarly: a one-in-ten chance of reversed roles more than doubles the set of acceptable programs. And a large common bond may earn a liquidity premium that a small one does not, so a facility kept deliberately small never becomes attractive enough to grow.

The implication is that the common budget rather than the common bond has to carry the safe country. A facility that only borrows and rebates asks its tax to do work it can barely do; one that also buys goods both members consume pays them directly and at a much lower rate.

1 Introduction

The European Union's fiscal response to the COVID-19 pandemic marked a turning point in its history, with the issuance of additional common debt in the order of €1 trillion. Political momentum has continued and the Draghi (2024) report called for a permanent common borrowing capacity to finance shared European priorities in defense, climate, and industrial policy. How to structure such a facility, and whether it can be made acceptable to all member states, has become one of the central challenges in European economic governance.

The difficulty is political as much as economic. The dominant view in both academic and policy debates treats fiscal integration as redistribution: fiscally strong core countries, such as Germany and the Netherlands, are reluctant to subsidize fiscally weak periphery countries, such as Italy and Greece. Under this framing, the political calculus for creditor countries is unfavorable by construction: any program that benefits borrowers must impose costs on lenders. This view shaped the rejection of Eurobonds after the 2010–2012 sovereign debt crisis and the insistence on strict conditionality in the NextGenerationEU architecture of 2020. It implies that the political obstacles to fiscal union reflect genuine economic trade-offs, not merely coordination failures.

This paper argues that this framing is incomplete. I show that a common fiscal program, in which a supranational institution issues common bonds funded through a non-defaultable common tax, can constitute a Pareto improvement: both the fiscally weak and the fiscally strong country can benefit, with no compensating transfer arranged outside the common instrument itself. The key institutional feature is the separation between common fiscal claims and national sovereign claims. The common tax remains enforceable when national debt is defaulted on and the institution can borrow at the safe rate. The risky country gains from a front-loaded safe-rate transaction it cannot arrange on its own account, while the same tax provides enough risk sharing to compensate the safe country. The gains are therefore gains from intertemporal trade and insurance rather than from a direct transfer.

I build on the sovereign-default framework of Dávila and Weymuller (2017), placing two countries, a risky periphery (R) and a safe core (S), in a currency area. A supranational institution, which I call the European Commission (EC), issues common bonds at the risk-free rate and distributes the proceeds symmetrically so that each country receives an amount θ. After this initial transfer, both countries borrow on their own national markets: the EC does not co-sign or guarantee any individual sovereign bond, and national debt is not mechanically retired. The program is funded by a proportional income tax τ levied on both countries, and the distinguishing feature is that the tax is non-defaultable: even if R defaults on its national bonds, the EC continues to collect the common tax. The policy problem is therefore two-dimensional, indexed by the pair (θ, τ).

This raises an immediate objection: if R can commit to a non-defaultable tax, why can it not commit to non-defaultable bond service, which would dissolve the default problem altogether?

Section 2.1 takes credibility to scale with the size of the obligation rather than to be all-or-nothing. The relevant size is net, not gross. The EC rebates its revenue to members, so what R would escape by repudiating the common claim is not the tax it pays but its share of the EC's interest cost. At the calibrated program that net obligation is well under a tenth of a percent of income, an order of magnitude below both the gross tax rate that compensates S and the national bond service R does default on. The issue is therefore quantitative: I locate the values of τ that deliver a Pareto improvement and ask whether the obligation they impose is small enough to be plausibly non-defaultable.

The fiscal arithmetic has a threshold. When the EC tax is low, the common budget can run a deficit in bad aggregate states, with tax revenues falling short of the facility's funding cost. A defaulting R is then shielded from its share of the EC deficit, reinforcing default incentives in the states where they are already strongest. I characterize the fiscal threshold $\bar{\tau}(\theta)$, the minimum tax rate above which this perverse incentive is fully eliminated in every income state. Above $\bar{\tau}(\theta)$, the EC always runs a surplus and R cannot gain from defaulting by escaping an EC deficit contribution.

The welfare effects of the program decompose into transparent channels for each country. When the transfer size θ expands, R receives a larger date-0 cash grant. Because the grant is unrestricted, any reduction in R's debt occurs through the country's equilibrium saving and issuance decisions. A larger program could in principle also reduce R's equilibrium default probability, raising the price lenders pay on its national bonds and benefiting R on its entire debt stock; the calibration finds that this credit response instead runs the other way. Non-defaultability is an enforcement property, making the EC's claim collectible in states where R repudiates its national debt, and it is distinct from any disciplining effect the program may or may not have on R's default decision.

S's welfare change has an analogous structure: it receives the same upfront transfer and bears the EC interest cost, but stands to benefit from a spillover channel, since any reduction in R's default frequency spares it the direct losses that arise whenever R defaults. The common tax also functions as a risk-sharing instrument: because the EC rebates revenues equally across members, τ channels resources toward whichever country draws the lower income realization each period. Because the levy and rebate are symmetric, these insurance transfers are zero-sum state by state and zero in expectation. Thus the tax shifts resources between members without favoring either one on average, and each pays only its share of the EC's interest cost in expectation.

At the calibrated parameters, I evaluate the introduction of a fresh common-debt program that raises the per-country transfer from $\theta = 0$ to $\theta = 0.1$ of annual output, a total issuance of about 10% of the combined annual GDP of the two countries. This program is Pareto improving only if common taxation is large enough, and the quantitative answer to *how* large is the main result. R benefits from the program at any τ because it receives a front-loaded grant and discounts the perpetual EC fee more heavily than the lenders who fund it. This is not a gain from cheaper national credit: the credit response works against R rather than for it. That gain is front-loaded

by construction, so the claim is about adoption rather than continuation: the program is a Pareto improvement to adopt, not one a member would choose to renew once the grant has been spent. S's welfare is negative at $\tau = 0$ and rises with τ, first as the fiscal exposure channel closes and then as risk sharing accumulates.

Two thresholds matter and they are an order of magnitude apart. The fiscal threshold $\bar{\tau}(0.1)$, above which the EC never runs a deficit, is a tax rate of approximately 0.09%: negligible, and therefore cheap to legislate. But clearing it only leaves S unharmed; it does not pay S for the program. Compensation arrives through the risk-sharing flow, which is proportional to τ, so the safe country turns strictly positive only at a rate of about 1.2%, thirteen times the fiscal threshold.

The cost the program imposes on the safe country is second-order throughout the grid, and a tax comparable in scale to the EU's current own-resources ceiling extinguishes it. The Pareto-improving region that results is real and contiguous, but for small values of τ it is confined to programs at or below roughly 12.5% of annual output, so the 10% benchmark sits close to its upper edge. The tax the program requires differs from current EU own resources in how it is collected rather than in its scale: the EC levies it directly instead of receiving it from member states. Nor does it only cover the facility's interest cost: it requires substantial revenue, which redistributes between members across income states.

For scale, the common debt already agreed at the EU level is on the order of €1 trillion, roughly 5% of EU GDP, though its distribution across member states is uneven and its duration is, at least in principle, limited.[1] A program of about 10% of GDP is therefore sizable but not unprecedented: common liabilities of a comparable magnitude already exist.

The analysis reframes the political constraint. A national politician cannot credibly support a program that makes their constituents worse off, regardless of aggregate European welfare gains. The question is therefore not whether the strong country should subsidize the weak one, but whether the institutional parameters (θ, τ) satisfy the condition that the safe country gains. This condition is checkable from observables: sovereign spreads, default frequencies, and cross-country output co-movement during crises, and is satisfied at empirically plausible parameter values. The institutional distance from the status quo to a Pareto-improving fiscal union is smaller than commonly assumed, though not in the dimension usually discussed: once the EC can levy the tax directly, the hard part is not enforcing the common claim, which is small, but paying the safe country enough to accept the program.

Section 5 considers three extensions. Public-good spending is the most important: by the standard fiscal-federalism logic, central provision of a jointly consumed good internalizes cross-border benefits and gives both countries a gain that does not have to be extracted from the other.

[1] Including the outstanding liabilities of the European Investment Bank, the European Stability Mechanism, and the European Financial Stability Facility brings the total closer to 10% of GDP, although the use of those funds is even more skewed across countries.

In the calibrated extension the safe country's welfare gain rises sharply and the Pareto-improving region expands from a narrow band confined to small programs to most of the policy grid. The binding constraint in the baseline is that S can be compensated only through a tax that must be large; the public good relaxes that constraint. It also reverses the ordering of the two thresholds: with a common public good S turns strictly positive at tax rates below the threshold above which the EC never runs a deficit, so paying the safe country now costs less than shielding it.

Which member occupies which fiscal role also need not be known when the arrangement is designed. The division of the euro area into a permanently risky periphery and a permanently safe core is a modeling convenience, and a country that expects to be safe may still weigh the possibility that the roles are allocated the other way. A one-in-ten chance of role reversal more than doubles the number of Pareto-improving policies and nearly doubles the largest program the safe country will accept.

Third, the rate the EC pays may itself depend on the scale of the program, which creates a coordination trap. Existing EU bonds have not achieved safe-asset status: a large issuance can earn a convenience yield while a small one pays an illiquidity premium, so S's welfare dips below the status quo at low θ and recovers only once the market is deep enough. Its participation set is consequently disconnected, and under a unanimity protocol that expands the program in small steps both no integration and full integration are stable with no path between them. Expectations about the depth of the common-bond market are self-fulfilling: a program kept small never earns the convenience yield that would justify enlarging it.

Related literature. This paper builds directly on the sovereign-default framework of Eaton and Gersovitz (1981), Aguiar and Gopinath (2006), Arellano (2008), and Chatterjee and Eyigungor (2012), as extended to a two-country setting by Dávila and Weymuller (2017). That paper varies pooled issuance and joint liability and finds that some joint-bond designs benefit both countries; here a common institution issues its own tax-backed debt without guaranteeing or retiring national debt. de Ferra and Mallucci (2025) are closest on political feasibility: a central authority that sets both members' borrowing lowers debt and default and raises joint welfare, but one member loses, which they read as making the arrangement politically infeasible. I leave national borrowing decentralized, give the institution its own claim and tax, and search directly for policies at which both members gain. Hatchondo et al. (2017) introduce capped non-defaultable borrowing in a single-country sovereign-default model; here the non-defaultable claim is issued by a common institution and backed by a cross-country tax.

The paper also speaks to work on joint European debt issuance. Pallara and Renne (2024) study how fiscal limits affect Eurobond pricing. Ando et al. (2023) are closest on the policy question but study a different object: their EDMA cancels legacy national debt and rolls its own debt over without a dedicated fiscal resource, so feasibility is a stochastic debt-sustainability condition. Jarociński and Maćkowiak (2018) also use a non-defaultable Eurobond, to resolve monetary-fiscal

indeterminacy rather than to study member-country participation. Bianchi et al. (2026) study a centralized Treasury that issues Eurobonds for countercyclical stabilization and coordinates with monetary policy, finding that both members benefit; gains come from reduced output and inflation volatility rather than an evaluation of country-specific adoption constraints. Tirole (2015) analyzes incentive-compatible bailouts, while Azzimonti and Quadrini (2024) show that cross-country default spillovers can make anticipated bailouts Pareto improving even before debt is chosen; this paper constructs a permanent ex-ante arrangement that requires no bailout.

A separate literature studies fiscal unions and risk sharing in currency areas. Farhi and Werning (2017) characterize optimal fiscal unions when nominal rigidities make private risk sharing constrained inefficient, Martinez et al. (2022) rank capital market union arrangements, and Aguiar et al. (2015) show that limited commitment in a monetary union generates overborrowing; the non-defaultable tax instead raises the cost of default without constraining issuance. Karaivanov et al. (2026) design incentive-compatible unemployment reinsurance without cross-country redistribution; here risk sharing is bundled with common-debt service. The risk-sharing literature also fixes the scale of the result: Lucas (1987) shows that insuring aggregate consumption is worth little in smooth developed economies, and Asdrubali et al. (1996) find that the US federal budget smooths only a modest share of state-specific shocks. On measurement I follow Dávila and Schaab (2025) in evaluating the reform from a predetermined initial state, but here welfare is reported country by country rather than aggregated, so the safe country's participation constraint is explicit.

The public-good extension is a fiscal-federalism exercise. Samuelson (1954) gives the efficiency condition, Oates (1999) and Besley and Coate (2003) frame the centralization trade-off, and Berriel et al. (2024) places that trade-off in a monetary union. I add sovereign default: the same EC tax that funds the common good is also the enforceable fiscal claim that changes national default incentives.

Finally, this paper connects to work on safe assets in Europe. Brunnermeier et al. (2017) propose European Safe Bonds, senior claims on a diversified portfolio of national sovereign bonds. Diversification and subordination create safety without joint liability or a common fiscal backstop, so the senior claim is safe by virtue of how existing default risk is repackaged. Leandro and Zettelmeyer (2019) and Cioffi et al. (2019) construct euro-area safe assets on the same principle, avoiding mutualization or redistribution by design. Pallara et al. (2025) reach the opposite conclusion, that jointly guaranteed eurobonds deliver their liquidity gains only under continuous issuance with the interest savings redistributed across members, and Bonfanti and Marcucci (2026) trace the failure of existing EU bonds to price as safe assets to the investors who hold them.

The rest of the paper is structured as follows. Section 2 develops the model and derives the fiscal threshold. Section 3 derives the welfare decompositions and characterizes the Pareto-improvement region analytically. Section 4 presents the calibration and the quantitative results. Section 5 discusses extensions, starting with public goods. Section 6 concludes.

2 A Model of Partial Fiscal Integration

This section develops a two-country sovereign-default model in which a common fiscal institution issues debt once and services it thereafter with a tax that no member can default on. Section 2.1 sets out the environment, the institution, its credibility, and the tax threshold above which the common budget is never in deficit. Section 2.2 states the budget constraints and value functions, and Section 2.3 the bond pricing and the equilibrium definition.

2.1 Environment and the Common Fiscal Institution

This section builds directly on the sovereign-default framework in Dávila and Weymuller (2017). Two countries, labeled R (risky) and S (safe), interact within a currency area as small open economies borrowing from competitive international capital markets populated by risk-neutral lenders. My contribution is to add a common fiscal institution that provides a one-time common-bond grant to both countries, funded through a non-defaultable common tax. I refer to this supranational institution as the European Commission (EC). Time is discrete and the horizon is infinite, $t = 0, 1, 2, \ldots$.

Each country is governed by a benevolent planner who maximizes the welfare of a representative consumer and receives a stochastic endowment each period. Log-income follows an AR(1) process:

$$\log y_{it} = \rho_i \log y_{i,t-1} + e_{it}, \tag{2.1}$$

where $\rho_i \in (0, 1)$ is the persistence parameter and the innovations are jointly normally distributed:

$$\begin{pmatrix} e_{Rt} \\ e_{St} \end{pmatrix} \sim \mathcal{N}(0,\, \Sigma)\,, \qquad \Sigma = \begin{pmatrix} \sigma_R^2 & \rho_{RS}\,\sigma_R\sigma_S \\ \rho_{RS}\,\sigma_R\sigma_S & \sigma_S^2 \end{pmatrix}. \tag{2.2}$$

The cross-country correlation ρ_{RS} generates positive income comovement, creating scope for risk sharing, while country-specific innovations generate dispersion. Denote by $y^{min} \equiv \min_{(y_R,\, y_S)}(y_R + y_S) > 0$ the minimum possible aggregate income over the discretized state space.

The representative consumer in country i has constant relative risk aversion (CRRA) preferences,

$$u_i(c) = \frac{c^{1-\gamma_i}}{1-\gamma_i}, \qquad \gamma_i \in (0,1) \cup (1, \infty), \tag{2.3}$$

and discounts future utility at rate $\beta_i \in (0, 1)$. In the calibration of Section 4.1 I set $\beta_S = q_S = 1/(1 + r_S)$, placing country S exactly at the knife-edge $\beta_S(1 + r_S) = 1$ rather than at the strict impatience $\beta_S(1 + r_S) < 1$ standard in small open economies, for the reason given in Section 3.2. Country S never defaults on its national debt; country R defaults strategically. Whenever R defaults, S suffers a fixed output loss $\varphi \geq 0$, a direct spillover standing in for the losses that

accompany a default event, as in Dávila and Weymuller (2017) and Tirole (2015); Azzimonti and Quadrini (2024) endogenize this loss through cross-border holdings of the defaulting sovereign's debt.

At date $t = 0$ the EC issues a perpetual bond at the rate $r_J = r_S$, so its issue price is $\bar{q} = 1/(1 + r_S) = q_S$. The bond carries a per-period service cost of $8(1 - \bar{q})\theta$, with the first payment falling due immediately at $t = 0$, and therefore raises $8(1 - \bar{q})\theta \sum_{t \geq 0} \bar{q}^t = 8\theta$ in cash. The proceeds are distributed symmetrically as a one-time lump-sum grant, giving each country 4θ in date-0 cash, and the EC retains on its own balance sheet a liability whose market value at issuance is 8θ, falling to $8\bar{q}\theta$ once the date-0 payment is made. Here and throughout, the program size θ is denominated in *annual* output. The model is quarterly and each country's output is normalized to one per quarter, so annual output is 4 and the program raises 4θ per country in model units; $\theta = 0.1$ is thus a transfer of ten percent of annual GDP.[2]

National debt stocks are not touched at $t = 0$: the common-bond liability stays with the EC. After $t = 0$ both countries continue to borrow at their own national rates: R at the endogenous market rate $q_R(b'_R, y)$ reflecting its default risk, and S at the risk-free rate q_S. The EC does not co-sign or guarantee any individual country's bonds; the two debt markets are fully segmented after the initial transfer.

In every period $t \geq 0$ the EC levies a proportional tax $\tau \geq 0$ directly on income in every member country to service the common bond, and distributes the net budget position equally between the two countries. The EC's period-t budget surplus is:

$$\omega_t = \tau(y_{R,t} + y_{S,t}) - 8(1 - \bar{q})\theta. \tag{2.4}$$

The first term is total tax revenue; the second is the total per-period debt-service cost of the common bonds. Each country receives (or contributes) one half of ω_t, so each country's net EC position is $\frac{1}{2}\omega_t = \frac{1}{2}\tau(y_{R,t} + y_{S,t}) - 4(1 - \bar{q})\theta$. The EC does not consume or invest: (2.4) is a mechanical accounting identity, and the budget position is immediately rebated to or collected from member countries.

Combining the levy with the rebate, country i's ongoing EC flow is $\frac{\tau}{2}(y_{j,t} - y_{i,t}) - 4(1 - \bar{q})\theta$. The insurance components sum to zero state by state and, because baseline incomes have equal unconditional means, each has zero expectation; only the common-debt fee remains on average.

The parameter θ is the per-country transfer and is the single instrument governing the extent of fiscal integration; all other features are fixed by symmetry or are policy parameters taken as given. In the baseline θ is exogenous and held constant, and the assumption $r_J = r_S$ is relaxed in Section 5.

[2]Placing the first service payment at $t = 0$ rather than at $t = 1$ makes the EC's budget (2.4) operate in every period $t \geq 0$. The choice is a pure relabeling: deferring it raises $8\bar{q}\theta$ rather than 8θ but also spares the period-0 fee, leaving the net date-0 receipt at $4\bar{q}\theta$ under either timing.

Non-defaultability of the EC tax. The tax τ is non-defaultable: even if R defaults on its national bonds, the EC continues to collect $\tau y_{R,t}$ from country R. Two features support this. The EC levies τ directly on the tax base, so no national treasury stands between the base and the center, and repudiating a treaty obligation could jeopardize a member's access to structural funds and the single market.

The tax is therefore an own resource in a narrower sense than EU budget law uses, since current own resources are transfers from national treasuries rather than levies the center collects.[3] A transfer of that kind is a claim on the sovereign. It corresponds to the deficit call $\frac{1}{2}|\omega_t|$, which a defaulting R can withhold, rather than to τ, which it cannot. The EU's own-resources ceiling thus benchmarks the scale of τ, not an existing collection power. Greece is a precedent for the priority members give EU claims rather than for direct collection: it restructured approximately €200 billion in national bonds in 2012 while remitting its EU budget contributions without interruption throughout.

The claim in question is also smaller than the tax rate suggests, because the EC rebates $\frac{1}{2}\omega_t$ to each member. Country i's net contribution is $\tau y_{i,t} - \frac{1}{2}\omega_t = \frac{1}{2}\tau(y_{i,t} - y_{j,t}) + 4(1-\bar{q})\theta$, where j is the other member: a zero-mean insurance transfer plus its share of the EC's interest cost. What R would escape by repudiating the common claim is therefore the fee $4(1-\bar{q})\theta$ and not the gross levy, a magnitude of the same order as the fiscal threshold derived below.

Appendix A.4 measures that obligation at the point where the temptation to walk away is greatest and shows it is small. Once the date-0 grant has been spent, a member faces only the perpetual fee and the ongoing tax, and Figure A.12 prices what it would gain by leaving at that point: between 0.06 and 0.08 percent of consumption at the reference program $\theta = 0.1$, and at most 0.35 percent at the largest program on the grid. The assumption is therefore load-bearing rather than vacuous, since that gain is strictly positive at every $\theta > 0$, but the claim the institution has to enforce is an order of magnitude smaller than the national bond service R does default on.

The asymmetry creates a *perverse fiscal incentive*. In a deficit state a repaying R contributes $\frac{1}{2}|\omega_t|$ to the EC while a defaulting R is shielded from the obligation, so default pays a fiscal reward on top of the debt relief it already delivers, and S absorbs the shortfall, which appears as the term $\min\{\frac{1}{2}\omega, 0\}$ in S's default-state budget constraint (2.9) below. Whether the EC can run a deficit at all is a question of arithmetic rather than of equilibrium, and the following threshold settles it.

Proposition 2.1 (Fiscal Threshold)**.** *Define the threshold tax rate:*

$$\bar{\tau}(\theta) \equiv \frac{8(1-\bar{q})\,\theta}{y^{min}}, \tag{2.5}$$

[3]Council Decision (EU, Euratom) 2020/2053, Articles 2, 3 and 9: member states make all categories of own resources available to the Commission, the GNI-based call is set each year to balance the budget, and own resources for payments are capped at 1.40% of EU GNI, raised temporarily by 0.6 percentage points until NextGenerationEU is repaid.

with y^{min} the minimum aggregate income defined in Section 2.1. Then:

(i) For any $\tau \geq \bar{\tau}(\theta)$: $\omega_t \geq 0$ for every realization (y_R, y_S) and every period t; the additional fiscal cost S bears due to R's default is exactly zero, S's total exposure to R's default is exactly φ, and the marginal fiscal incentive for R to default, the gain from escaping a deficit contribution, is identically zero.

(ii) Any increase $\Delta\theta > 0$, accompanied by an adjustment in τ to maintain $\tau \geq \bar{\tau}(\theta + \Delta\theta)$, leaves S's fiscal exposure to R's default unchanged at φ.

Proof. See Appendix Section A.1. □

Proposition 2.1 is a direct implication of the EC budget constraint (2.4) and requires no equilibrium argument. It holds regardless of preferences, stochastic structure, or parameter values.[4] Clearing $\bar{\tau}(\theta)$ is not sufficient for the safe country to accept the program. It returns S to indifference about the EC's budget position; it does not pay S anything. Section 3.3 takes up what does.

Why a non-defaultable tax but defaultable bonds? A natural objection is that if R can commit to a non-defaultable tax, it should be able to commit equally to non-defaultable bond service. I take credibility to scale with the size of the claim rather than to be all-or-nothing: an obligation small relative to what a centrally collected levy can raise, and to the treaty penalties non-compliance triggers, is enforceable, while one large enough to rival the gains from walking away is not. National bond service is large by construction, which is what makes default tempting; the EC tax is credible only insofar as it stays small.

I therefore ask the reverse question rather than assume an arbitrary τ is feasible: how large would a non-defaultable tax have to be to deliver a Pareto improvement, and is that magnitude small enough to be credible? Section 4.2 answers it. The region is reached at a tax rate on the order of 1% at $\theta = 0.1$, an order of magnitude above the fiscal threshold but comparable in scale to the EU's current own-resources ceiling. The obligation that rate imposes is the net fee $4(1 - \bar{q})\theta$ rather than the gross levy.

2.2 Budget Constraints and Value Functions

This subsection writes down the resource constraints of the two countries, the cost of default and the terms on which R returns to the market, and the recursive problems each country solves. Let $d_R \in \{0, 1\}$ denote R's default decision in the current period, and suppress the time subscript. The period budget constraints are as follows.

[4] Its one dependence is on the discretization: y^{min} is the smallest income on the grid, so a finer grid extends the lower tail and raises $\bar{\tau}(\theta)$ slightly.

When R repays ($d_R = 0$):

$$c_R = (1-\tau)y_R - b_R + q_R(b'_R, y)\, b'_R + \tfrac{1}{2}\omega, \tag{2.6}$$

$$c_S = (1-\tau)y_S - b_S + q_S b'_S + \tfrac{1}{2}\omega. \tag{2.7}$$

When R defaults ($d_R = 1$):

$$c_R = (1-\tau)y_R - L(y_R) + \max\{\tfrac{1}{2}\omega,\ 0\}, \tag{2.8}$$

$$c_S = (1-\tau)y_S - \varphi - b_S + q_S b'_S + \tfrac{1}{2}\omega + \min\{\tfrac{1}{2}\omega,\ 0\}. \tag{2.9}$$

Following the convention in the sovereign default literature, debt is denominated as a positive liability stock: b_i denotes debt due this period and b'_i next-period issuance. The risky country's issuance is restricted to $b'_R \in [-\underline{a}, \bar{b}_R]$ for a borrowing limit $\bar{b}_R > 0$ and a small savings allowance $\underline{a} > 0$; the allowance plays no quantitative role, since R's ergodic mass on the savings nodes is exactly zero at the calibrated parameters, but it is what places the model inside the uniqueness condition discussed after Definition 2.1.

The safe country's issuance is restricted to a band $b'_S \in [\underline{b}_S, \bar{b}_S]$ with $0 < \underline{b}_S < \bar{b}_S$, because the knife-edge $\beta_S = q_S$ makes S's debt a driftless random walk, which is stationary only between two reflecting barriers; the pair pins S's ergodic debt-to-output ratio, and neither barrier carries an economic interpretation of its own.[5]

The budget constraints (2.6)–(2.9) govern each country's decisions in every period $t \geq 0$, the EC surplus ω_t included; at $t = 0$ each country's resources additionally carry the 4θ of cash from the EC lump-sum grant. Recording that grant as a reduction in the date-0 liability is without loss of generality. Because date-0 issuance b'_i is chosen freely, cash in hand and a lower inherited debt are the same object: grouping the date-0 resources in (2.6) as $-b_i + 4\theta = -(b_i - 4\theta)$, the problem of a country holding debt b_i and receiving a grant 4θ is *identical* to that of a country holding debt $b_i - 4\theta$ and receiving nothing.

The split between immediate consumption and deleveraging is therefore set by the country's own policy function rather than by the EC, and the two countries resolve it differently: S, nearly indifferent to the timing of consumption at the knife-edge, absorbs almost the whole transfer into a permanently lower debt path, while the impatient R consumes the larger part. Section 3.1 evaluates welfare at that shifted liability.

The safe-country constraint (2.9) carries the two channels through which R's default harms S. The term $\min\{\frac{1}{2}\omega, 0\} \leq 0$ is the *fiscal channel*: a defaulting R still claims the rebate $\frac{1}{2}\omega$ when the EC

[5]Some such device is unavoidable, since a small open economy at or near $\beta_S(1+r_S) = 1$ has no stationary debt distribution, and Schmitt-Grohé and Uribe (2003) show that other standard remedies deliver nearly identical dynamics at standard calibrations. The band is chosen among them because it is the one device that leaves the knife-edge itself undisturbed.

is in surplus but cannot be compelled to cover its share of a deficit, so when $\tau(y_R + y_S) < 8(1-\bar{q})\theta$ the uncovered shortfall $4(1-\bar{q})\theta - \frac{1}{2}\tau(y_R + y_S)$ falls on S on top of its own contribution. The term $\varphi \geq 0$ is the *direct spillover channel* of Section 2.1, incurred regardless of the income state.

Country R defaults strategically each period by choosing $d_R \in \{0, 1\}$. Defaulting eliminates R's obligation to repay its national debt but imposes a direct output cost, in the tradition of Eaton and Gersovitz (1981). Following Chatterjee and Eyigungor (2012) and Dávila and Weymuller (2017), I use a reduced-form default loss that is zero in bad income states and convex in good ones:

$$L(y_R) = \min\Big\{\max\big\{0,\ d_0\, y_R + d_1\, y_R^2\big\}\,,\ y_R\Big\}, \qquad d_0 < 0 < d_1, \tag{2.10}$$

so the loss is zero below the onset income $\hat{y}_R = -d_0/d_1$ and rises convexly above it; the defaulting country consumes $y_R - L(y_R)$ net of the tax and EC transfer. The outer cap at y_R makes default-state consumption non-negative by construction and never binds over the ergodic range at the calibrated values. Default is therefore most attractive in bad income states, where the cost vanishes, and least attractive in good ones, where the convex loss bites. Following a default, country R is excluded from bond markets and regains market access each period with probability $\alpha > 0$.

Partial recovery at re-entry. Defaulted debt is not written off in full. Creditors recover a fraction $1 - h$ of the defaulted face, where $h \in [0, 1]$ is the haircut, and R re-enters the market carrying the reinstated obligation

$$b_R^{\text{re}} \;=\; (1-h)\, b_R, \tag{2.11}$$

where b_R is the stock outstanding at the time of default. Setting $h = 1$ recovers the zero-recovery benchmark of Arellano (2008), Chatterjee and Eyigungor (2012) and Dávila and Weymuller (2017), in which R re-enters debt-free.

Risky country. Country R's value function solves:

$$V_R(b_R, y) = \max_{d_R \in \{0,1\}} \Big\{(1 - d_R)\, V_R^c(b_R, y) + d_R\, V_R^d(b_R, y)\Big\}\,, \tag{2.12}$$

where $y \equiv (y_R, y_S)$ is the vector of current income realizations. The value of repaying is:

$$V_R^c(b_R, y) = \max_{b'_R \in [-\underline{a}, \bar{b}_R]} \big\{u_R(c_R) + \beta_R\, \mathbb{E}\big[V_R(b'_R, y')\big]\big\}\,, \tag{2.13}$$

$$\text{subject to (2.6).}$$

The value of defaulting depends on the defaulted stock as well as on income, because the stock determines what R carries back into the market at re-entry:

$$V_R^d(b_R, y) = u_R\big((1-\tau)y_R - L(y_R) + \max\{\tfrac{1}{2}\omega,\, 0\}\big) \\ + \beta_R\Big\{\alpha\,\mathbb{E}\big[V_R\big((1-h)b_R,\, y'\big)\big] + (1-\alpha)\,\mathbb{E}\Big[V_R^d(b_R, y')\Big]\Big\}. \tag{2.14}$$

Since the income process (2.1) follows a first-order Markov chain, the expectation $\mathbb{E}[\cdot \mid y]$ is conditional on the current income state, so the bond price $q_R(b'_R, y)$ depends on the current state as well as on next-period debt and the whole problem is solved on the full (b_R, y) state space.

Safe country. Country S does not hold R's bonds, so R's debt enters S's problem only through whether R has market access. I therefore let S condition on a regime indicator $m \in \{\mathsf{c}, \mathsf{d}\}$, recording whether R is in the market (c) or excluded (d), rather than on R's debt stock itself. Country S takes as given the law of motion of that indicator, which R's solution determines: the probability of moving from c to d is R's default probability conditional on current income, evaluated at R's ergodic debt distribution,

$$\lambda(y) \;\equiv\; \Pr\big(m' = \mathsf{d} \mid m = \mathsf{c},\, y\big) \;=\; \int d_R(b_R, y)\, d\mu_R^{\mathsf{c}}(b_R \mid y), \qquad \Pr\big(m' = \mathsf{c} \mid m = \mathsf{d}\big) = \alpha, \tag{2.15}$$

where μ_R^{c} is the ergodic distribution of R's debt over the in-market pool. Country S's state is then (b_S, y, m), and it carries one value function per regime. When R is in the market, S chooses b'_S before knowing whether R repays this period, so both its current consumption and its next-period regime are contingent on that draw. Write c_S^{n} for consumption in (2.7), the branch in which R repays, and c_S^{d} for consumption in (2.9), the branch in which R defaults and S pays the spillover and any uncovered EC deficit; both are evaluated at the same b'_S. Then

$$V_S^c(b_S, y) = \max_{b'_S \in [\underline{b}_S, \bar{b}_S]} \Big\{\big(1-\lambda(y)\big)u_S(c_S^{\mathsf{n}}) + \lambda(y)\, u_S\Big(c_S^{\mathsf{d}}\Big) \\ + \beta_S\Big[\big(1-\lambda(y)\big)\mathbb{E}\, V_S^c(b'_S, y') + \lambda(y)\,\mathbb{E}\, V_S^d(b'_S, y')\Big]\Big\}. \tag{2.16}$$

The default probability entering (2.16) is $\lambda(y)$, conditional on *current* income, as (2.15) defines it: R's decision this period is taken at this period's income, and it is what moves the regime from c to d. When R is out of the market,

$$V_S^d(b_S, y) = \max_{b'_S \in [\underline{b}_S, \bar{b}_S]} \Big\{u_S(c_S) \\ + \beta_S\,\mathbb{E}\Big[\alpha\, V_S^c(b'_S, y') + (1-\alpha)\, V_S^d(b'_S, y')\Big]\Big\}, \\ \text{subject to (2.9).} \tag{2.17}$$

The continuation value in (2.17) accounts for the possibility that R remains excluded next period with probability $1-\alpha$. Country S therefore bears the spillover in the period of the default itself, through the $\lambda(y)$ branch of (2.16), and again in every period of the exclusion spell that follows. While R is excluded, S continues to borrow and save freely and incurs the spillover cost φ; the EC program remains in place and ω_t is computed using R's realized income even during exclusion.

Conditioning on m rather than on b_R is a deliberate simplification, and it is an approximation rather than a relabeling. The exact problem would have S carry R's debt as a state and forecast default from it, which requires S to track R's issuance policy as well as its default rule. Because R's debt affects S only through the indicator d_R and through the aggregate income entering ω_t, replacing the state-contingent default probability by its in-market ergodic average is accurate wherever R's debt distribution is tight, as Appendix Figure A.3 shows it to be here, and it removes R's debt grid from S's state space while preserving the interpretation of the spillover channel.

2.3 Bond Pricing and Equilibrium

This subsection prices the two countries' bonds, shows how partial recovery enters the risky price schedule, and defines the recursive equilibrium. Risk-neutral lenders earn zero expected profits, so the safe country's bonds and the EC's common bonds, neither of which carries default risk, are priced at the risk-free rate:

$$q_S = \bar{q} = \frac{1}{1+r_S}. \tag{2.18}$$

For the risky country, lenders price in both the probability of default and what they recover if it occurs. With persistent income, next period's default probability depends on R's debt choice and the current income state, so the price schedule is

$$q_R(b'_R, y) \;=\; \frac{\mathbb{E}[\Psi(b'_R, y') \mid y]}{1+r_R}, \qquad \Psi(b,y) \;=\; \big(1-d_R(b,y)\big) \;+\; d_R(b,y)\,\Phi(b,y), \tag{2.19}$$

where $\Psi(b,y)$ is the value to a lender of one unit of face maturing in state (b,y): it is repaid in full if R repays, and is worth the defaulted claim $\Phi(b,y)$ otherwise. The value of that defaulted claim is itself forward-looking, because the lender is paid only once R regains market access and begins servicing the reinstated stock (2.11). Per unit of *original* face,

$$\Phi(b,y) \;=\; \frac{1}{1+r_R}\,\mathbb{E}\big[\,\alpha\,(1-h)\,\Psi\big((1-h)b,\,y'\big) \;+\; (1-\alpha)\,\Phi(b,y') \;\big|\; y\big]\,. \tag{2.20}$$

With probability α the sovereign re-enters and the claim converts into $1-h$ units of face on the reinstated stock, worth Ψ each; with probability $1-\alpha$ exclusion continues and the claim retains its value Φ. The factor $1-h$ multiplying Ψ is the face conversion. Together, (2.19)–(2.20) make q_R the present value of exactly the payments R makes, including those it makes after re-entry. At

$h = 1$ we have $\Phi \equiv 0$ and $\Psi \equiv 1 - d_R$, and (2.19) collapses to the familiar zero-recovery expression $q_R = (1 - \mathbb{E}[d_R])/(1 + r_R)$.

The discount rate $r_R \geq r_S$ carries no convenience yield, and the wedge $r_R - r_S \geq 0$ is not double-counted with the default premium. Safe-country bonds provide liquidity and store-of-value services beyond their pecuniary return, so lenders accept a lower pecuniary yield on S's debt, as in He et al. (2019) and Dávila and Weymuller (2017); R's bonds provide no such services, so q_R in (2.19) compensates lenders for default risk alone and R borrows at that price throughout, as it appears in (2.6).

The two claims on R are therefore priced on different margins. The EC's claim is non-defaultable, so it is priced at $\bar{q}$ and never enters (2.19), while national issuance moves R down its own schedule, so the marginal revenue it raises, $q_R + b'_R \, \partial q_R / \partial b'_R$, falls short of the average price; with one-period debt this operates within the period rather than across maturities. Resources delivered through the EC do not reprice the debt R already owes, which is why R cannot replicate the transfer on its own account.

Definition 2.1 (Recursive Equilibrium)**.** *A recursive equilibrium consists of the value functions*

$$\{V_R, V_S, V_R^c, V_R^d, V_S^c, V_S^d\}$$

and the policy functions

$$\{b'_R(\cdot), b'_S(\cdot), d_R(\cdot)\}.$$

It also includes a defaulted-claim value Φ *and a risky bond price function*

$$q_R : [-\underline{a}, \bar{b}_R] \times \mathcal{Y} \longrightarrow \big[0, \, 1/(1 + r_R)\big].$$

These objects satisfy:

(i) V_R, V_R^c, V_R^d *satisfy* (2.12)–(2.14)*, with* b'_R *and* d_R *the associated optimal policies.*

(ii) V_S^c, V_S^d *satisfy* (2.16)–(2.17) *given the regime chain* (2.15) *implied by* $d_R(\cdot)$ *and* μ_R^{c}*, with* b'_S *the associated optimal policy.*

(iii) The bond price function satisfies (2.19)*, with* Φ *solving* (2.20)*, for all admissible* $b' \in [-\underline{a}, \bar{b}_R]$ *and all income states* $y \in \mathcal{Y}$*.*

The equilibrium is unique for each (θ, τ), which the normative analysis needs: it compares welfare across a grid of policies, and differences could not be attributed to the policy if the solution at a given policy were not pinned down. Auclert and Rognlie (2016) restore uniqueness in this one-period-bond benchmark under *either* a strictly positive bound on asset holdings *or* a default punishment that is strictly costly in every income state; the savings allowance $\underline{a} > 0$ delivers the first

outright. Partial recovery (2.11) points the same way, since a defaulting sovereign carries $(1-h)b_R$ back into the market even in the low-income states where the output loss (2.10) vanishes. Nothing below rests on it.[6]

The equilibrium has a block-recursive structure: R's problem is solved first, extracting the default rule d_R and the ergodic distribution behind (2.15). Country S's problem is then solved taking that regime chain as given. The separation runs one way because S does not hold R's bonds, R's bond price depends only on R's own choices, and both countries face the same EC program parameters (τ, θ), which are taken as fixed; nothing S does feeds back into R's problem.

3 Welfare Decompositions and Pareto Improvements

This section derives analytical expressions for how the welfare of each country changes as the fiscal integration parameters (θ, τ) vary. Section 3.1 introduces the welfare measures and the stochastic discount factor notation used throughout. Sections 3.2 and 3.3 derive the marginal welfare decompositions with respect to θ and τ, each closing with a discussion of what the corresponding instrument must deliver for a Pareto improvement; the formal conditions are collected in Appendix A.2.

The decompositions below are first-order objects and are read as such, in two senses. Each states how $\mathcal{W}_i$ responds to a *marginal* change in one instrument at a given policy, so the channels are attributions valid in a neighborhood of that policy rather than an accounting identity for a finite reform. Each is also a *partial* rather than a total derivative: it collects the flow effects along the equilibrium path together with the equilibrium responses of q_R and of R's default frequency, but holds the invariant distribution fixed. Section 4.3 measures the residual that integrating them along a finite path leaves.

3.1 Welfare Measures and Stochastic Discount Factors

This subsection fixes the welfare measure and the notation the decompositions use: lifetime utility evaluated at the pre-reform distribution, reported in consumption-equivalent units, and summed along equilibrium paths with each country's own stochastic discount factor.

Hold the equilibrium of Section 2 fixed at a policy pair (θ, τ) and define, for country $i \in \{R, S\}$ starting from state (b_i, y), the value function

$$V_i(b_i, y;\, \theta, \tau) \;\equiv\; \mathbb{E}_{(b_i,y)}\left[\sum_{t=0}^{\infty} \beta_i^t\, u_i(c_{it})\right], \tag{3.1}$$

[6] Their proof is for i.i.d. shocks whereas income here is persistent, so uniqueness is supported by their argument rather than implied by it. I verify it numerically by their bracketing check: iterating the price map from the minimal and from the maximal default threshold reaches the same fixed point.

with the expectation taken over equilibrium paths initialized at that state. Being the discounted sum of all future flow utilities from the inherited state, V_i already integrates the entire equilibrium response to the program, including the gradual transition of the cross-sectional distribution toward its new invariant law.

Let $\mu^*(\,\cdot\,;\theta,\tau)$ denote the stationary distribution induced by the equilibrium policy functions at (θ,τ). The program is a once-and-for-all reform, announced and enacted at date 0 with the economy at the pre-reform distribution $\mu^*(\,\cdot\,;0,0)$, so the welfare-relevant object is the lifetime utility of the agents alive at that date: the new-regime value function averaged over the *pre-reform* distribution,

$$\mathcal{W}_i(\theta,\tau) \;\equiv\; \int V_i(b_i - 4\theta,\; y;\,\theta,\tau)\, d\mu^*(b_i, y;\,0,0). \tag{3.2}$$

The shifted argument is the date-0 grant: by the argument of Section 2.2, receiving 4θ in cash and inheriting a liability lower by 4θ are the same problem, so the grant requires no separate term.[7] This object is *transition-inclusive*: it credits the program with the entire path of welfare effects as the distribution adjusts, because the forward-looking $V_i(\cdot;\theta,\tau)$ already prices that adjustment. It is the discrete counterpart of the welfare assessment of Dávila and Schaab (2025), who differentiate lifetime utility from a predetermined initial state and treat steady-state gains as the long-run limit of a full transition path.

The measure is kept country by country rather than aggregated into a single planner's objective. Aggregation would require taking a stand on the relative weight placed on each member state, which is precisely the judgment a Pareto criterion avoids and the one that national governments will not delegate. The regime is permanent, so the bond-price schedule, the default policy and the value functions jump to their new-regime values at once, and only the cross-sectional distribution transitions.[8]

Welfare changes are reported in consumption-equivalent units rather than in utils. Under CRRA preferences with $\gamma_i \in (0,1)\cup(1,\infty)$ lifetime utility scales as $V_i((1+\kappa)c) = (1+\kappa)^{1-\gamma_i} V_i(c)$, so the permanent proportional transfer κ_i that leaves country i indifferent between the no-integration

[7] Two qualifications, both concerning states in which R lacks market access and therefore cannot convert cash into a debt path. First, the program is assumed to be initiated only in states in which R is in the market, so (3.2) is evaluated against the date-0 distribution conditional on that event, renormalized, for both countries jointly: the reported object is the welfare gain conditional on initiation in normal times, and the probability of that event is itself higher under larger programs. Second, whether R defaults *in* the initial period is a decision taken after the transfer arrives rather than a state that can be conditioned upon, so it is retained: the relevant date-0 value for R is $\max\{V_R^c(b_R - 4\theta, y),\, V_R^d(b_R, y) + u_R(c_R^d(y) + 4\theta) - u_R(c_R^d(y))\}$, with c_R^d given by (2.8).

[8] Appendix A.4 reports two alternatives and keeps them apart: the same reform launched from each policy's own invariant distribution, which leaves the sign of the Pareto region unaffected, and the conventional stationary comparison, which drops the grant and is adverse.

baseline and the equilibrium at (θ, τ) solves $(1+\kappa_i)^{1-\gamma_i} = \mathcal{W}_i(\theta,\tau)/\mathcal{W}_i(0,0)$:

$$\kappa_i(\theta,\tau) \;=\; \left[\frac{\mathcal{W}_i(\theta,\tau)}{\mathcal{W}_i(0,0)}\right]^{1/(1-\gamma_i)} - 1. \tag{3.3}$$

Three scales. The paper reports welfare on three scales, which are not comparable to one another. *Utils* are the level $\mathcal{W}_i$, negative throughout under CRRA, so only differences across policies are interpretable; because κ_i is a strictly increasing transform of the utility difference, the two always agree in sign. *Consumption equivalents* are κ_i from (3.3), written $\Delta\mathcal{W}_i(\theta,\tau)$ and quoted in percent, and are what every heatmap and every welfare number in Section 4 reports. *Date-0 consumption units* are the normalized marginal objects $u_i'(c_{i0})^{-1} d\mathcal{W}_i/d\zeta$ of (3.5), in which the four decompositions and Table 4.2 are expressed; they are per unit of the instrument and are never comparable with the consumption-equivalent percentages.

Stochastic discount factors. For each country i, define the *stochastic discount factor* (SDF) along the equilibrium path:

$$\Pi_{it} \;\equiv\; \frac{\beta_i^t\, u_i'(c_{it})}{u_i'(c_{i0})}. \tag{3.4}$$

By construction $\Pi_{i0} = 1$, and a unit change in c_{it} is worth Π_{it} units of date-0 consumption to country i. The normalized welfare change from a shift $d\zeta$ in a generic policy instrument $\zeta \in \{\theta, \tau\}$ is therefore

$$\frac{1}{u_i'(c_{i0})}\,\frac{d\mathcal{W}_i}{d\zeta} \;=\; \mathbb{E}\left[\sum_{t=0}^{\infty} \Pi_{it}\,\frac{\partial c_{it}}{\partial \zeta}\right], \tag{3.5}$$

where the partial derivative holds the household's *debt policy functions* fixed (envelope theorem), and the expectation is over the ergodic distribution of equilibrium paths. Dividing by $u_i'(c_{i0})$ puts (3.5) and the four decompositions that follow from it in date-0 consumption units, the third of the three scales above.[9]

Partition of equilibrium paths. The consumption change $\partial c_{it}/\partial\zeta$ depends on whether R is in the bond market, so I partition equilibrium paths into two events: $\mathcal{N}$, in which R is in the market and repays ($d_{Rt} = 0$), and $\mathcal{D}$, in which it is out ($d_{Rt} = 1$), whether because it defaults in that period or because it remains excluded following an earlier default. The two partition every date. For any

[9] (3.5) and its specializations (3.8)–(3.15) decompose flow effects along the equilibrium path, augmented by selected equilibrium-response terms, rather than exact total derivatives of ergodic welfare. The envelope theorem eliminates first-order effects of the household's own debt policy, but q_R, R's default frequency and the measure of the event sets are not its choice and are reintroduced where relevant. An exact total derivative would also require perturbing the invariant distribution μ^*; Section 4.2 evaluates the magnitudes numerically instead.

adapted process X_t and either event $\mathcal{E}$, define the *cumulative path-weighted expectation*

$$E^{\mathcal{E}}[X_t] \;\equiv\; \mathbb{E}\left[\sum_{t=0}^{\infty} \mathbf{1}_{\{\mathcal{E} \text{ at } t\}}\, X_t\right], \tag{3.6}$$

with $E^{\text{all}} \equiv E^{\mathcal{N}} + E^{\mathcal{D}}$. Both sum from $t = 0$, because the EC's tax, rebate and debt-service operations run in every period from the issuance period onwards, as in Section 2.1. The date-0 grant is the one item outside them: it enters every decomposition below as a separate deterministic term, contributing 4 per unit of θ to date-0 resources.

Working above the fiscal threshold. The four decompositions below are derived under $\tau \geq \bar{\tau}(\theta)$, so that $\omega_t \geq 0$ in every state by Proposition 2.1. Below the threshold the EC can run a deficit, and since a defaulting R is shielded from its share of it while S absorbs the whole shortfall, default periods must then be split by the sign of ω_t, which adds one threshold-contingent term to each decomposition. Appendix A.3 carries that general case. The restriction costs little because it restricts the *policy*, not the equilibrium, and because every Pareto-improving policy in the calibrated economy turns out to lie above the threshold.

3.2 Welfare Effects of Common Debt Increases

Common debt is the instrument that pays R. It moves resources to date 0, where the impatient country values them most, and its cost is a per-period fee that both countries carry for as long as the bond is outstanding.

Differentiate the budget constraints with respect to θ, holding the household policy functions fixed. Since $\omega_t = \tau(y_{Rt} + y_{St}) - 8(1 - \bar{q})\theta$, we have $\partial\omega_t/\partial\theta = -8(1 - \bar{q})$.

Risky country. The only date-0 channel is the lump-sum grant. In (3.2) it enters as the state shift $b_R - 4\theta$ rather than as a term in date-0 consumption, but the two coincide at the margin: differentiating gives 4 times the marginal value of a unit of date-0 resources, which at the optimum is $u'_R(c_{R0})$, so after the normalization in (3.5) the channel contributes 4 per unit of θ. For $t \geq 0$ in repayment states, $c_{Rt} = (1 - \tau)y_{Rt} - b_{Rt} + q_R(b'_{Rt}, y_t)\, b'_{Rt} + \frac{1}{2}\omega_t$. Since R borrows entirely at the market price q_R, the envelope theorem gives:

$$\left.\frac{\partial c_{Rt}}{\partial \theta}\right|_{\mathcal{N},\, t\geq 0} \;=\; \underbrace{b'_{Rt}\,\frac{\partial q_R}{\partial \theta}}_{\text{credit response}} \;\underbrace{-4(1 - \bar{q})}_{\text{EC interest cost}}, \tag{3.7}$$

where b'_{Rt} is next-period issuance. The effect of θ through the *optimal debt choice* vanishes at the optimum, but $\partial q_R/\partial\theta$ survives: it is an *equilibrium* price response, and q_R enters the budget constraint directly.

I name this term the *credit response* because it captures whether a common debt program induces R to default less and hence pay a lower interest rate in equilibrium. No assumptions guarantee a positive sign, and the calibration of Section 4 finds it negative for small programs and effectively zero for $\theta \geq 0.175$.

In default states $c_{Rt} = (1-\tau)y_{Rt} - L(y_{Rt}) + \frac{1}{2}\omega_t$, the max in (2.8) being slack above the threshold, so $\partial c_{Rt}/\partial\theta = -4(1-\bar{q})$ there as well: R pays its share of the fee in every period, in the market or out of it. Collecting all terms:

$$\frac{1}{u'_R(c_{R0})}\frac{d\mathcal{W}_R}{d\theta} = \underbrace{4}_{\text{date-0 cash grant}} - \underbrace{4(1-\bar{q})\,E^{\text{all}}[\Pi_{Rt}]}_{\text{EC interest cost}} + \underbrace{E^{\mathcal{N}}\left[\Pi_{Rt}\, b'_{Rt}\,\frac{\partial q_R}{\partial\theta}\right]}_{\text{credit response (full debt)}}. \tag{3.8}$$

The EC interest cost is unambiguously negative. Netted against the grant, however, the pair is strictly positive for R, and the reason is an impatience arbitrage rather than any gain from insurance or improved credit. The EC borrows at r_S, so the upfront cash is 4 per unit of θ, while R discounts the perpetual fee $4(1-\bar{q})$ at a rate governed by $\beta_R \ll \bar{q}$: approximating $E^{\text{all}}[\Pi_{Rt}] \approx 1/(1-\beta_R)$, as Section 4.3 does, the present value of the fees is about $4(1-\bar{q})/(1-\beta_R)$, far below 4 whenever R is sufficiently impatient. The approximation replaces the ratio $u'_R(c_{Rt})/u'_R(c_{R0})$ by its ergodic average of one; it is not an upper bound.

The credit-response channel would be positive if $\partial q_R/\partial\theta > 0$, that is if a larger program lowered R's equilibrium default probability and raised the price on its *entire* national debt. Its sign here is not favorable, and the reason is structural. The scale of the program enters the recursive equilibrium through exactly one term, the per-period funding cost $-8(1-\bar{q})\theta$ in (2.4); the date-0 grant is a one-time transfer that shifts the initial state without altering any policy function or price schedule. At fixed τ, a larger θ is therefore a pure levy, and above the threshold R bears it whether it repays or not. Default discipline in this model is bought with τ, not with θ: the tax cushions R's income in bad states, while the role of θ is to supply the front-loaded resources and to make that tax non-defaultable.

Safe country. At $t = 0$, S also receives the date-0 cash grant, worth 4 per unit of θ. In every period S borrows at the fixed rate q_S, so $\partial q_S/\partial\theta = 0$ and the only ongoing effect operates through ω_t. Above the threshold the fiscal-exposure term $\min\{\frac{1}{2}\omega_t, 0\}$ in (2.9) is zero in every state, so S's consumption carries $\frac{1}{2}\omega_t$ and nothing else from the EC:

$$\frac{\partial c_{St}}{\partial\theta} = \tfrac{1}{2}\cdot(-8(1-\bar{q})) = -4(1-\bar{q}). \tag{3.9}$$

Beyond the mechanical EC cost, S is also affected by *behavioral* changes in R's default probability: to the extent that a larger θ leads R to default less often, a response that is an equilibrium outcome rather than a property of the policy and is assumed in condition (d) of Proposition A.1, S avoids the spillover cost φ it incurs in $\mathcal{D}$ states.

To isolate this behavioral channel, let $\Delta_t \equiv -\partial \Pr_t(d_R = 1 \mid y_t)/\partial\theta$ denote the induced *change* in period-t default frequency *conditional on the income state*, read as a local comparative static on that conditional default rate rather than as a pointwise derivative of the discrete default indicator. Conditioning on income rather than aggregating to a single number is what lets Δ_t be paired with a state-varying exposure in Appendix A.3, where a defaulting R costs S its share of the EC deficit as well as φ: that extra exposure is a function of ω_t and hence of aggregate income alone, so it is constant on the set over which Δ_t is averaged and comes out of the conditional expectation exactly.[10] No sign is imposed on Δ_t: it is positive when a larger program disciplines R and negative when it does not, and Section 4.2 finds the latter. Collecting all terms:

$$\frac{1}{u_S'(c_{S0})}\frac{d\mathcal{W}_S}{d\theta} \;=\; \underbrace{4}_{\text{date-0 cash grant}} \;-\; \underbrace{4(1-\bar{q})E^{\text{all}}[\Pi_{St}]}_{\text{EC interest cost}} \;+\; \underbrace{\varphi\, E^{\text{all}}[\Pi_{St}\,\Delta_t]}_{\text{spillover channel}}. \tag{3.10}$$

S receives the same date-0 grant as R and, unlike R, cannot waive the per-period fee $4(1-\bar{q})$ in any state. Under the calibration of β_S adopted in Section 4.1 these two terms cancel *exactly* along any interior equilibrium path, rather than merely approximately.

Proposition 3.1 (Cash-flow neutrality of the EC program for S)**.** *Suppose $\beta_S = q_S = 1/(1+r_S)$ and assume that along the equilibrium path neither barrier of the band $[\underline{b}_S, \bar{b}_S]$ binds with positive probability.*[11] *Then the date-0 cash grant and the present value of the perpetual EC fee cancel exactly:*

$$4 \;=\; 4(1-\bar{q})\sum_{t=0}^{\infty}\mathbb{E}_0[\Pi_{St}]. \tag{3.11}$$

Proof. See Appendix Section A.1. □

The proposition isolates the *normative channel* through which the program reaches S: absent the spillover channel, S would be completely indifferent to the magnitude of θ, so any net gain or loss above the threshold is attributable to that one term. One natural implementation of that

[10] Replacing the state-level default response by its income-conditional average is a further approximation, exact only under a conditional-independence condition that fails here; Appendix A.2 states why, and the implementation carries the same approximation. It is a property of the presentation rather than of the argument: Proposition A.1, where the sign of this channel does any work, signs the state-level integrand directly and never invokes it.

[11] The two barriers cut in opposite directions when they bind, and the remark following the proof in Appendix Section A.1 signs each. Only the deleveraging limit $\underline{b}_S$ operates in the calibration, and it works against S; footnote 16 measures it.

indifference is for S to use the grant to reduce risk-free borrowing by 4θ and hold the lower debt path thereafter, so that the extra national-debt service exactly offsets the EC fee period by period. The model does not impose it: the grant is an unrestricted transfer and, at the knife-edge, S is indifferent across timing paths with the same present value, so (3.11) is a present-value identity rather than a prediction for the debt path. Keeping S's Euler equation interior in every state is the analytical reason for adopting the knife-edge $\beta_S = q_S$.[12]

That surviving term inherits the sign of Δ_t and is therefore not signed by theory: it is a benefit only if a larger program disciplines R. Section 4.2 finds that it does not, so in the calibrated economy the channel is a cost and θ makes S worse off.

The safe country is the binding side. Along θ the constraint $d\mathcal{W}_R \geq 0$ asks nothing of the equilibrium, since R's gain is secured by the grant alone. Everything therefore turns on S, whose (3.10) reduces to the sign of the spillover channel. Appendix A.2 states sufficient conditions for S's marginal welfare attribution to be non-negative along a policy path in θ, which is what the decomposition can deliver and is one step short of a statement about $\mathcal{W}_S$ itself: two of the four restrict the policy path and can be imposed by construction, while the other two, a non-negative risk-sharing flow and a default region shrinking in θ, are properties of the equilibrium. The calibration satisfies the first and rejects the second, so θ on its own does not carry S. What remains is the other instrument, to which I now turn.

3.3 Welfare Effects of Common Taxation Increases

The common tax is the instrument that pays S. Above the fiscal threshold it operates as pure risk sharing, moving resources from whichever country is currently richer to whichever is poorer, and it is the only channel through which S is compensated for the program.

Differentiate the budget constraints with respect to τ. Since $\partial\omega_t/\partial\tau = y_{Rt} + y_{St}$, the direct effect runs through two channels, the tax on own income and the rebate on aggregate income, both applying from $t = 0$ as in (3.6).

Risky country. In repayment states, $c_{Rt} = (1-\tau)y_{Rt} - b_{Rt} + q_R(b'_{Rt}, y_t)\, b'_{Rt} + \frac{1}{2}\omega_t$, so differentiating with respect to τ and applying the envelope theorem to b'_{Rt} (while retaining the direct price response $\partial q_R/\partial\tau$ that is not eliminated by envelope):

$$\left.\frac{\partial c_{Rt}}{\partial\tau}\right|_{\mathcal{N}} = -y_{Rt} + \tfrac{1}{2}(y_{Rt} + y_{St}) + b'_{Rt}\,\frac{\partial q_R}{\partial\tau} = \frac{y_{St} - y_{Rt}}{2} + b'_{Rt}\,\frac{\partial q_R}{\partial\tau}. \tag{3.12}$$

[12] Setting $\beta_S = q_S$ places S exactly at the Aiyagari–Bewley impatience boundary $\beta_S(1+r_S) = 1$ of Aiyagari (1994), so S's marginal utility is a martingale in interior states and its stationary debt distribution is not pinned down by impatience but only by the location of the borrowing limit.

The tax is a *risk-sharing device*: R pays in proportion to its own income and receives back half the joint revenue, so the net transfer is positive exactly when $y_{St} > y_{Rt}$, that is in good times for S and bad times for R. A second channel mirrors the θ decomposition: if τ disciplines default, q_R rises on R's entire outstanding stock, a credit-response gain proportional to $b'_{Rt}\,\partial q_R/\partial\tau$.

In default states the credit-response term is absent, since R issues no bonds, but the risk-sharing algebra is unchanged, because $\max\{\frac{1}{2}\omega_t, 0\} = \frac{1}{2}\omega_t$ above the threshold. Collecting:

$$\frac{1}{u'_R(c_{R0})}\,\frac{d\mathcal{W}_R}{d\tau} \;=\; \underbrace{E^{\text{all}}\left[\Pi_{Rt}\,\frac{y_{St}-y_{Rt}}{2}\right]}_{\text{risk-sharing}} \;+\; \underbrace{E^{\mathcal{N}}\left[\Pi_{Rt}\,b'_{Rt}\,\frac{\partial q_R}{\partial\tau}\right]}_{\text{credit response}}. \tag{3.13}$$

The risk-sharing channel is positive when R is poor relative to S in high-marginal-utility states, the canonical gain from fiscal transfers. The credit-response channel is the analog of its counterpart in (3.8).

Safe country. Above the threshold the fiscal-exposure term is inactive in every state, so

$$\frac{\partial c_{St}}{\partial\tau} = -y_{St} + \tfrac{1}{2}(y_{Rt}+y_{St}) = \frac{y_{Rt}-y_{St}}{2}. \tag{3.14}$$

The two direct resource effects are equal and opposite, so they sum to zero state by state. This physical symmetry does not imply welfare neutrality: each country values transfers across states with its own stochastic discount factor, allowing both risk-sharing terms to be positive. The only other route to S is the spillover, through the response of R's default frequency to the tax, $\Delta^\tau_t \equiv -\partial\Pr_t(d_R=1)/\partial\tau$, defined as Δ_t was in (3.10):

$$\frac{1}{u'_S(c_{S0})}\,\frac{d\mathcal{W}_S}{d\tau} \;=\; \underbrace{E^{\text{all}}\left[\Pi_{St}\,\frac{y_{Rt}-y_{St}}{2}\right]}_{\text{risk-sharing}} \;+\; \underbrace{\varphi\,E^{\text{all}}[\Pi_{St}\,\Delta^\tau_t]}_{\text{spillover channel}}. \tag{3.15}$$

S's risk-sharing channel is the mirror image of R's. The two countries share an unconditional income distribution, but the SDF weighting breaks the symmetry, since the states in which S's marginal utility is high need not be those in which $y_{Rt} > y_{St}$. The sign therefore depends on the joint distribution of (y_{Rt}, y_{St}) and the SDFs, and I determine it numerically in Section 4.2.

Behavioral effects of τ. The spillover term in (3.15) is the only route by which R's response to the tax reaches S above the threshold, and theory does not sign it: it is a benefit if τ lowers R's default frequency and a cost otherwise. Table 4.2 measures it at -0.003 against a risk-sharing term of 1.985, so it does not change the reading of (3.15).

Necessary conditions for a Pareto improvement. Above the threshold the task of compensating S falls entirely to risk sharing and to the behavioral response just described. Proposition A.2 in Appendix A.2 records what this requires. Above $\bar{\tau}$, a marginal increase in τ raises both countries' welfare attributions only if R's risk-sharing term, augmented by the credit-response channel, is non-negative, and only if S's risk-sharing term, whose sign is calibration-dependent, covers any loss on its spillover channel through φ. They are *necessary* where the θ-side conditions are sufficient, and they rest on the sign of the risk-sharing terms rather than on an assumption about the default rule. Like the θ-side conditions, they are stated for the attributions of footnote 9 rather than for the exact total derivatives, so they are necessary for a Pareto improvement itself only up to the residual that Section 4.3 measures. Section 4 measures both sets of conditions.

4 Quantitative Analysis

This section takes the decompositions of Section 3 to a calibrated economy. Section 4.1 sets the parameters and reports the fit at the no-integration benchmark. Section 4.2 maps the welfare surfaces of both countries over the policy grid and locates the Pareto-improving region. Section 4.3 evaluates each channel of the decompositions at the calibrated equilibrium and identifies which of them drive the result.

4.1 Calibration

Parameter values. The exercise assigns to the standard parameters of the sovereign default model the values conventional in the quarterly literature, and calibrates only the few that the specific question requires. Both countries have the same constant relative risk aversion $\gamma = 2$, in line with the sovereign default literature following Arellano (2008) and Dávila and Weymuller (2017). The risk-free rate differential $r_R - r_S = 0.1\%$ captures the non-pecuniary convenience yield on safe assets: lenders accept a lower pecuniary return on S's bonds because they provide liquidity and store-of-value services not offered by R's bonds. In annual real terms r_R and r_S equal about 1.0% and 0.6%, matching the average real yields on Italian and German sovereign debt over a sample beginning in 2002. The safe country's discount factor is set to $\beta_S = q_S = 1/(1 + r_S)$, so that $\beta_S(1 + r_S) = 1$ exactly and S is at the knife-edge between net saving and borrowing in steady state. The haircut on defaulted debt is set to $h = 0.5$, so lenders recover half of the defaulted face and R re-enters carrying $(1 - h)b_R$ as in (2.11). This sits in the middle of the range the empirical literature reports, between Yue (2010)'s 0.27 and the 0.63 implied by Cruces and Trebesch (2013)'s mean net-present-value haircut.

The default cost function (2.10) follows Chatterjee and Eyigungor (2012) and Dávila and Weymuller (2017), with $d_0 = -0.59 < 0 < d_1 = 0.84$, so the implied loss ratio $L(y_R)/y_R = d_0 + d_1 y_R$ is

25% of output at mean income.[13] The re-entry probability $\alpha = 0.05$ implies an expected exclusion duration of $1/\alpha = 20$ quarters, or five years, following a default event, in line with empirical estimates of market re-access. The fixed spillover cost that S bears in each period during which R is in default is set to $\varphi = 0.02$, a per-period consumption loss of about two percent while R is excluded.[14] The income process (2.1) is a bivariate AR(1) in logs with persistence $\rho_R = \rho_S = 0.95$ and innovation standard deviation $\sigma_R = \sigma_S = 0.03$, close to the quarterly values in Dávila and Weymuller (2017). The cross-country innovation correlation is $\rho_{RS} = 0.50$, generating substantial scope for risk sharing through the common tax. The EC bond rate is set to the safe country's risk-free rate, $r_J = r_S = 0.15\%$, consistent with the baseline assumption $\bar{q} = q_S$ of Section 2. Section 5 relaxes this by introducing an illiquidity premium on EC bonds. Table 4.1 summarizes all parameter values.

Table 4.1: Baseline parameter values

The table reports the parameter values of the baseline calibration. The model period is a quarter, so the three interest rates are quarterly, and all of them are real, worth about 1.0% (r_R) and 0.6% ($r_S = r_J$) a year. All values are set externally to estimates standard in the quarterly sovereign default literature, except β_R, d_0 and d_1, which are calibrated jointly to the four benchmark moments reported below. Debt grids and the discretization of the income process are set out in Appendix A.7.

Symbol	Description	Value
$\gamma_R = \gamma_S$	CRRA risk aversion	2.0
β_R	Discount factor, risky country (R)	0.93
β_S	Discount factor, safe country (S)	$q_S \approx 0.9985$
r_R	Lender discount rate for R (no convenience yield)	0.25%
r_S	Risk-free rate, S (with convenience yield premium)	0.15%
r_J	EC bond rate	0.15%
φ	Fixed spillover cost (S incurs when R defaults)	0.02
d_0	Default cost, linear coefficient, Eq. (2.10)	-0.59
d_1	Default cost, quadratic coefficient, Eq. (2.10)	0.84
h	Haircut on defaulted debt (recovery $1-h$), Eq. (2.11)	0.50
α	Re-entry probability (20 quarters = 5 years exclusion)	0.05
$\rho_R = \rho_S$	AR(1) persistence of log income	0.95
$\sigma_R = \sigma_S$	Innovation std. dev. (quarterly)	0.03
ρ_{RS}	Cross-country innovation correlation	0.50

Benchmark fit. At the no-integration benchmark $(\theta, \tau) = (0, 0)$ the calibration targets moments from a representative sample of European countries in the period 2002–2025 and the model returns

[13] This is needed to fit a debt ratio near 105% of annual GDP with one-period debt: absent long maturity, the model needs a large output loss to sustain that stock at the observed default frequency.

[14] Empirical estimates of cross-border default spillovers vary widely, but φ enters the safe country's problem alone, so it shapes the location of S's participation margin and leaves R's gains unaffected. Appendix Figure A.6 brackets the baseline with $\varphi \in \{0, 0.05\}$ and shows that it shapes the Pareto region only at its margin.

the following moments: $E[b_S/y_S] = 61.5\%$, $E[b_R/y_R] = 104.5\%$, a default frequency of 1.09% per year, and an annualized spread over the safe rate of 1.02%. Appendix Figures A.3, A.4 and A.5 report the ergodic objects behind these moments: the debt-to-output distributions of both countries, the distribution of R's spread, and R's equilibrium bond-price schedule.

4.2 Welfare Surfaces and the Pareto Region

This subsection maps the two welfare surfaces over the policy grid, locates the cells at which both countries gain, and reports the tax rate at which the safe country is compensated. The two surfaces are two orders of magnitude apart in size: what the program does to S is second-order in either direction, and the work of the tax is to remove a cost rather than to deliver a large gain. The Pareto-improving region that follows is real and contiguous but confined to small programs paired with tax rates an order of magnitude above the fiscal threshold.

I compute $\mathcal{W}_R(\theta, \tau)$ and $\mathcal{W}_S(\theta, \tau)$ over a grid of (θ, τ) pairs, in each case solving for the recursive equilibrium of Section 2 and integrating against the ergodic distribution as described in Appendix A.7. Welfare is measured relative to the baseline with no common program, $(\theta, \tau) = (0, 0)$.

Figure 4.1: Welfare changes relative to no integration

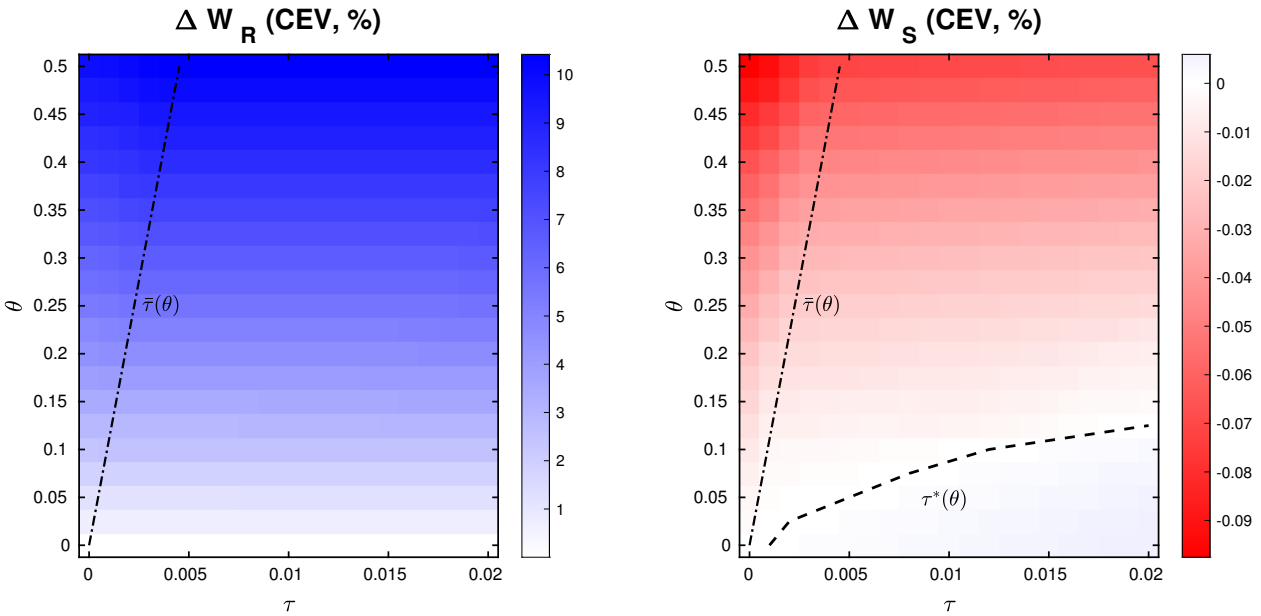

The figure shows the consumption-equivalent welfare change of each country, computed as in (3.3), at every point of the (θ, τ) policy grid and measured against the no-integration benchmark $(\theta, \tau) = (0, 0)$. *Left*: risky country R. *Right*: safe country S. Darker shading indicates larger gains (in blue) or losses (in red), on color scales set separately for the two panels, whose ranges differ by roughly two orders of magnitude. The dash-dotted curve on both panels is the fiscal threshold $\bar{\tau}(\theta)$ of Proposition 2.1. The dashed curve on the right panel is the compensating tax $\tau^*(\theta)$, the smallest grid tax at which both countries gain, which traces the lower-tax edge of the Pareto-improving cells. The grid is 21×21, evenly spaced over $\theta \in [0, 0.5]$ and $\tau \in [0, 0.02]$, so each surface is drawn from 441 solved recursive equilibria.

$\mathcal{W}_R$ is increasing in θ across the entire grid, consistent with the prediction of (3.8), reaching a

welfare gain of 10.4% of consumption at the upper end of the θ grid.[15] Variation in R's welfare is driven almost entirely by θ: the tax rate moves it only marginally, and Section 4.3 identifies which channel is responsible.

Figures A.7 and A.8 in the appendix show that θ does not discipline R. Holding $\tau = 0$ and raising θ from 0 to the top of the grid makes R's credit *worse*, not better: the annual default frequency rises from 1.09% to 1.15%, the spread widens from 1.02% to 1.06%, and R's ergodic debt ratio falls from 104.5% to 99.8% as it deleverages against a stiffer price schedule. A common-debt program that transfers resources without collecting taxes is a pure fiscal expansion from R's point of view, and it is the associated EC deficit, borne disproportionately in bad states, that erodes the price of R's national debt. Taxation reverses all three: at the largest program on the grid, moving τ from 0 to 0.005 restores the default frequency to 1.11%, the spread to 1.03%, and the debt ratio to 103.3%.

The safe country's welfare surface $\mathcal{W}_S$ is roughly two orders of magnitude smaller in absolute terms, confined to the range -0.098 to $+0.006$ percent against 0 to $+10.4$ percent for R.[16] The asymmetry of that range is the central finding of this subsection: S's losses from an untaxed program are more than sixteen times larger than the largest gain taxation can deliver to it anywhere on the grid, so what the tax buys S is the removal of a second-order cost and not a first-order gain. Issuing common debt without common taxation is negative for S, since the program then acts as an unconditional transfer financed jointly by both countries and the spillover channel does not compensate for it.

An increase in taxation eventually reverses the sign, but lying to the right of $\bar{\tau}(\theta)$ is neither necessary nor sufficient for it: at $\theta = 0.3$, 18 of the 21 tax-grid cells clear the threshold, yet $\Delta\mathcal{W}_S$ remains around -0.025 throughout that slice. The cells where S gains are confined to the low-θ, high-τ corner. Reaching them requires far more taxation than merely clearing the threshold. Appendix Figure A.6 shows that these results are robust to the calibrated value of φ.

The size of S's gain is what the international risk-sharing literature predicts. Lucas (1987) bounds the value of eliminating aggregate consumption fluctuations at roughly one tenth of one

[15] The gain is the date-0 grant net of the present value of the EC fee, worth $4 - 4(1 - \bar{q})/(1 - \beta_R)$ per unit of θ, as Section 4.3 shows. It is therefore proportional to the wedge between R's discount factor and the rate at which the EC funds itself, which is the same wedge the sovereign default literature relies on to generate borrowing at all, and it collapses only as $\beta_R \to \bar{q}$. Its magnitude is insensitive to where in that literature's range β_R is set: 3.940 at $\beta_R = 0.90$, 3.914 at the calibrated 0.93, 3.873 at Arellano (2008)'s 0.953, and 3.800 at 0.97, a spread of under two percent.

[16] Part of S's measured loss is the reflecting band of Section 2 rather than the program: cash-flow neutrality requires S to absorb the grant as a permanently better debt position, and the barrier $\underline{b}_S$ caps how far that position can improve. Re-solving S's problem against a band more than twice as wide, holding R's equilibrium fixed at $\tau = 0.02$, recovers between 0.37 and 0.47 of the measured $\Delta\mathcal{W}_S$ at three program sizes spanning the grid, the share falling as θ rises. These are point re-solves, so they scale the artifact without relocating the frontier. The band thus accounts for roughly two-fifths of the measured change, and it works in one direction only, biasing the Pareto region toward being too small.

percent of consumption. Gains of hundredths of a percent are the normal magnitude for a risk-sharing arrangement between two developed economies with smooth, persistent income, and the common tax delivers considerably less than complete markets. The empirical counterpart is Asdrubali et al. (1996), who find that the US federal tax-and-transfer system smooths about 13 percent of state-specific income shocks, an operative channel that is nonetheless the smallest of the three they measure. The euro area has no comparable supranational stabilizer: Ferrari and Rogantini Picco (2023) estimate negligible smoothing through international transfers among the original eleven members, with no significant change attributable to monetary union.

The Pareto region requires non-negligible taxation. Taxation extinguishes the cost to S at a rate comparable in scale to the EU's current own-resources ceiling, and where it does both countries gain: combining the two surfaces, 78 of the 441 grid cells are Pareto improving, and every one of them lies at $\theta \leq 0.125$.[17] That ceiling on θ is set by the tax grid rather than by economics: $\mathcal{W}_S$ is still rising in τ at $\tau_{\max} = 0.02$ at every θ, so the region has no interior upper edge in the tax dimension and $\theta \leq 0.125$ describes the region explored at tax rates up to 2% rather than the region that exists.[18] A Pareto improvement therefore exists, over a contiguous region, without side payments. What that set costs, however, is that the tax rates involved are not small.

The program sizes are in the range the policy literature contemplates, though the two bounds are different objects. Ando et al. (2023) calculate that an EDMA with no dedicated fiscal resource could issue about 15% of euro-area GDP beyond NextGenerationEU by capitalizing the convenience yield on European safe assets; the bound here is instead the largest program satisfying both countries' participation constraints when national debt remains defaultable. The proximity is therefore not an external validation, but it does place the welfare-feasible scale in the range current policy work treats as institutionally relevant.

The dashed line on the right panel of Figure 4.1 makes the gap visible. It is the *compensating tax* $\tau^*(\theta)$, the smallest grid tax at which both countries are strictly better off than at $(0, 0)$, so it traces the lower-tax edge of the Pareto-improving cells and stops at the largest program the grid supports. It rises steeply, reaching 1.2% at $\theta = 0.1$ and the top of the explored range at $\theta = 0.125$. Read against the dash-dotted $\bar{\tau}(\theta)$ on the same panel, which reaches only 0.090% at $\theta = 0.1$, the compensating tax runs nine to eighteen times the fiscal threshold across program sizes.

Proposition 2.1 guarantees that clearing $\bar{\tau}$ eliminates S's excess fiscal exposure and R's perverse default incentive, but that only returns S to indifference about the deficit; it does nothing to *pay* S

[17] $\Delta\mathcal{W}_S$ carries numerical noise of order 0.001 percentage points, small in absolute terms but a sixth of S's entire positive range and far larger than the value of the single Pareto-improving cell at $\theta = 0.125$, which is itself only 0.000002. The interior of the region is robust; the exact location of its frontier, and the $\theta = 0.125$ row in particular, should not be read as precise. The band bias of footnote 16 is larger than this noise and runs in one direction only, so the region is if anything understated.

[18] Extrapolating the measured slope of $\mathcal{W}_S$ in τ, between 0.29 and 0.45 percentage points per unit, a grid reaching $\tau = 0.04$ would extend the region to about $\theta = 0.20$.

for the program. Compensation comes from the risk-sharing flow $\frac{1}{2}\tau(y_R - y_S)$, which is proportional to τ and therefore requires the EC to run a substantial surplus, not merely a balanced budget. The program cannot be sold to the safe country as a budget-neutral facility that covers its own interest: it has to be a tax with real revenue attached, redistributing between members across income states without favoring either one in expectation, and large enough that the insurance it provides outweighs the transfer embedded in the date-0 grant. Its 1.2% magnitude at $\theta = 0.1$ lies below the current 1.4%-of-GNI ceiling for payment appropriations, but this is a comparison of fiscal scale only, not of collection technology.

4.3 Quantifying the Channels

The decompositions of Sections 3.2–3.3 are not merely qualitative: every term can be evaluated at the calibrated equilibrium, and the relative magnitudes show which forces actually drive the welfare results. I compute each channel in (3.8), (3.10), (3.13), and (3.15) at every grid point using the ergodic-stationary approximation

$$E^{\mathcal{E}}[\Pi_{it}\, X_t] \;\approx\; \frac{1}{1-\beta_i} \sum_{s \in \mathcal{E}} \mu(s)\, \frac{u_i'(c_i(s))}{\bar{u}_i'}\, X(s), \tag{4.1}$$

where μ is the ergodic distribution over states s and $\bar{u}_i'$ the ergodic-average marginal utility. Table 4.2 reports the result at the calibrated transfer $\theta = 0.1$ and a tax rate $\tau = 0.01$ comfortably above the fiscal threshold. These are first-order objects evaluated at a stationary approximation, so they attribute the forces at work rather than account for them exactly, and Section 3 reads them accordingly.[19]

Three facts emerge from Panel A. First, the date-0 cash grant net of the EC interest cost is the only large positive contribution to R's gain from a larger program: that channel is worth 3.914 per unit of θ, and it is flat to four decimals across all 441 grid cells, because under (4.1) it equals $4 - 4(1-\bar{q})/(1-\beta_R)$ exactly. R receives the face value of the program at date 0 and discounts the perpetual fee that services it at $\beta_R = 0.93$, so it keeps almost all of the transfer.

This is an impatience arbitrage, not a gain from insurance or from improved credit. What leaves it unexploited is the slope of R's own price schedule: issuing nationally moves R down it, so the marginal revenue $q_R + b_R'\, \partial q_R / \partial b_R'$ falls short of $\bar{q}$ at the optimum, while the EC's claim raises the same resources without entering (2.19). At the benchmark R's average bond price is 0.996, a quarter of a percent below $\bar{q}$, while its mass-weighted marginal revenue is about 0.91, some nine percent below it. The level of R's spread is therefore not what blocks the transaction; the slope of

[19] Integrated over the finite policy change to $(\theta, \tau) = (0.1,\, 0.02)$, the channels recover about three-fifths of the exact welfare difference for R. The residual is concentrated in the credit-response term, a windowed slope of a price schedule that moves discretely with R's default set, so the error is largest in the transition band $\theta \in (0,\, 0.175)$; the other channels are identities or are flat in θ.

Table 4.2: Welfare channels at $\theta = 0.1$, evaluated above the fiscal threshold ($\tau = 0.01$)

The table reports the channels of the welfare decompositions of Section 3, evaluated at the single policy cell $(\theta, \tau) = (0.1, 0.01)$, a transfer of ten percent of annual output paired with a tax rate well above the fiscal threshold $\bar{\tau}(0.1) = 0.00090$. Panel A decomposes $d\mathcal{W}_i/d\theta$ into the channels of (3.8) and (3.10), Panel B decomposes $d\mathcal{W}_i/d\tau$ into those of (3.13) and (3.15). Entries are marginal contributions in date-0 consumption units per unit of the instrument, the scale of (3.5), and are not comparable with the consumption-equivalent percentages of Figure 4.1. Each is evaluated at the stationary approximation (4.1), so the entries attribute the forces at work at the margin rather than account exactly for a finite reform. † Default–EC tension for R, fiscal-exposure reduction for S; both vanish above $\bar{\tau}$. ‡ $\partial q_R/\partial\tau$ for R; $-\partial \Pr(d_R{=}1)/\partial\tau$ for S. The corresponding Panel A entries are $\partial q_R/\partial\theta$ and $\varphi\,\Delta_t$.

	Country R	Country S
Panel A: $d\mathcal{W}_i/d\theta$ (per unit of θ)		
Date-0 cash grant − EC interest cost	3.914	0.000
Fiscal exposure	—	0.000
Credit response (R) / spillover (S)	−2.259	−0.010
Total $d\mathcal{W}_i/d\theta$	1.656	−0.010
Panel B: $d\mathcal{W}_i/d\tau$ (per unit of τ)		
Risk-sharing	0.071	1.985
Threshold-contingent†	0.000	0.000
Credit response (R) / spillover (S)‡	0.000	−0.003
Total $d\mathcal{W}_i/d\tau$	0.071	1.982

the schedule is.

Second, R's credit response is substantially *adverse* at small programs: −2.259 at this cell, −3.853 at $\theta = 0$, and effectively zero for $\theta \geq 0.175$, so it is exactly where the Pareto region lies that it eats into the grant, offsetting 98 percent of it at $\theta = 0$ and 58 percent at the reference cell. It is also erratic there, moving between those magnitudes and near zero across adjacent θ nodes, which is the transition band the footnote above flags. This is the marginal counterpart of the comparative static reported above: enlarging the program worsens R's bond price, so θ is not a disciplining device.

Third, for S the cash grant and the present value of the EC fee cancel to machine precision, the exact numerical counterpart of the cash-flow-neutrality result of Proposition 3.1: the entry is 0.000 not by rounding but as an identity that holds cell by cell across the entire policy grid.

The three facts together sharpen the paper's central claim. Above the fiscal threshold (3.10) has only two terms besides the spillover channel and they cancel exactly, by Proposition 3.1. Enlarging θ on its own does essentially nothing for S, so the entire scope for making S strictly better off, and hence the entire Pareto-improving region, is created by τ. Common debt must be paired with common taxation.

Panel B shows where S's gains come from. Risk sharing dominates both countries' marginal welfare: the effects $d\mathcal{W}_R/d\tau$ (0.071) and $d\mathcal{W}_S/d\tau$ (1.982) are each accounted for almost entirely

by the risk-sharing term, with the threshold-contingent channels equal to zero above $\bar{\tau}$ and the equilibrium-response channels negligible. The risk-sharing gain to S is roughly *twenty-eight times* that to R. The asymmetry is a patience effect: S values a permanent flow at $1/(1-\beta_S) \approx 668$ against R's $1/(1-\beta_R) \approx 14$. That ratio of 47 overstates the measured 28, the gap reflecting the periods R spends excluded and the different states in which the two countries' marginal utilities are high.

Figure 4.2: The θ channels above the fiscal threshold, along the $\tau = 0.01$ slice

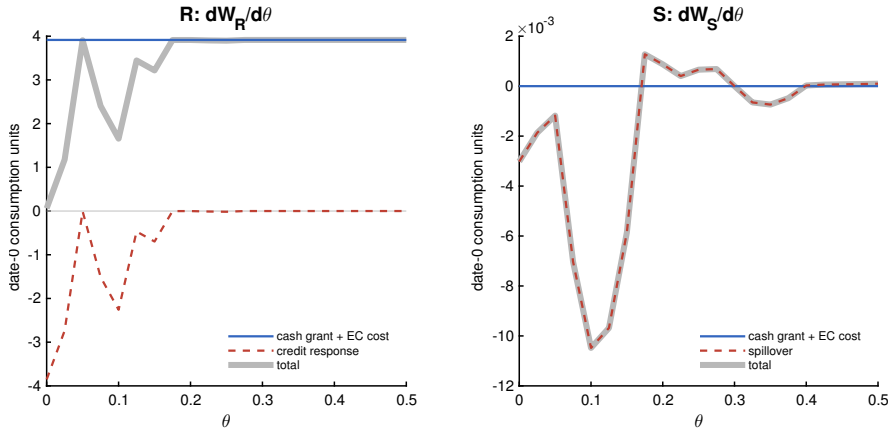

The figure shows how the channels of the θ decompositions vary with the scale of the program, along the slice of the policy grid at $\tau = 0.01$, which lies above $\bar{\tau}(\theta)$ for every θ on the grid. *Left*: the channels of (3.8) for the risky country. *Right*: those of (3.10) for the safe country. Table 4.2 reports the $\theta = 0.1$ point of both panels, in the same units. The total is the pale band and the channels that sum to it are drawn on top, so a channel lying on the band accounts for the total on its own; that is the case in the right panel, where the grant and the fee cancel and S's spillover channel *is* $d\mathcal{W}_S/d\theta$. S's threshold-contingent channel is identically zero on this slice and is not drawn. The τ decompositions are not plotted, because above the threshold each is its risk-sharing channel up to a negligible equilibrium-response term. Appendix Figures A.9 and A.10 show every channel of both decompositions over the full grid.

Figure 4.2 extends the table to the whole θ dimension at the same tax rate, and adds the one thing a single cell cannot show: the credit response decays unevenly, because it is the local slope of a price schedule that moves with discrete changes in R's default set. The remaining channels vary little by comparison, so S's panel stays close to zero across the entire range rather than lying exactly on it.

Below the threshold two further channels switch on, derived in Appendix A.3: S's fiscal-exposure term in (A.9) turns negative as it absorbs R's unpaid deficit share, and R's default–EC tension in (A.10) turns negative because a higher tax in deficit-default states buys R no rebate. Both collapse to zero exactly at $\bar{\tau}$, so the channel decomposition reproduces the fiscal threshold from the welfare side. Figures A.9 and A.10 in the appendix plot the full (θ, τ) profile of every channel.

5 Extensions: Additional Sources of Welfare Gains

Three of the baseline's assumptions shape the size of the Pareto region, and relaxing each of them widens it. The first is that the common budget only issues and rebates: Section 5.1 lets it also buy a jointly consumed public good. The second is that fiscal roles are known when the program is designed: Section 5.2 makes the assignment uncertain. The third is that common bonds pay the safe-country rate: Section 5.3 lets that rate be a function of the amount issued.

5.1 Common Public Goods and the Fiscal Commons

A common budget that also buys something both members want is Pareto improving over almost the entire policy grid, whereas the purely financial facility of Section 4 is not. This subsection lets the EC devote part of its surplus to a jointly consumed public good such as defense, cross-border energy and transport infrastructure, or joint research. Doing so raises the safe country's gain by two orders of magnitude and reverses the ordering of the paper's two tax thresholds.

Preferences and the common budget. I augment the baseline utility by a log public-good term:

$$U_i \;=\; E_0\left[\sum_{t=0}^{\infty}\beta_i^t\big(u_i(c_{it})+\vartheta\log G_t\big)\right], \qquad i\in\{R,S\}, \tag{5.1}$$

where u_i is the CRRA felicity defined in (2.3), $G_t>0$ is a flow public good consumed equally by both countries, and $\vartheta=0.1$ governs the weight on public consumption.[20]

The EC spends $\xi\geq 0$ units of the consumption good per period on the common good, which enters its budget as a direct cost and modifies the surplus (2.4) to

$$\omega_t(\xi) \;=\; \tau(y_{Rt}+y_{St})-8(1-\bar{q})\theta-\xi. \tag{5.2}$$

Each country continues to receive $\frac{1}{2}\omega_t(\xi)$, so central provision reduces each country's per-period transfer by $\xi/2$; the budget constraints (2.6)–(2.9) are unchanged in form. The countries also provide the good themselves and do so without coordinating, so total provision is $G_t=g_{Rt}+g_{St}+\xi$ with each $g_{it}\geq 0$ chosen taking the partner's contribution as given. Appendix A.5 develops the local provision game in full.

Underprovision and the case for centralizing. Uncoordinated national provision is inefficiently low, and that gap is what centralized provision exists to close. A contributing country equates the *private* marginal benefit to its own marginal utility, $\vartheta/G_t^N=u_i'(c_{it}^N)$, while a planner

[20] The value $\vartheta=0.1$ targets an aggregate public-good expenditure share of roughly 5–10% of GDP, consistent with a broad interpretation of jointly consumed public goods.

internalizing both countries' benefits sets $2\vartheta/G_t^{SP} = u_i'(c_{it}^{SP})$. The Nash level is too low because each country ignores the half of the benefit that accrues abroad. Centralization closes part of that gap only because at least one country sits at the corner $g_{it} = 0$, where it cannot absorb the levy funding ξ by cutting a contribution it is not making.

That is the generic configuration, since a single G_t^N cannot satisfy both first-order conditions at unequal marginal utilities.[21] This is the classic fiscal-federalism trade-off between internalizing spillovers and imposing a uniform policy (Oates, 1999), and the present application is a favorable one: defense, cross-border infrastructure and joint research have large spillovers, and the EC budget provides insurance rather than a one-way subsidy, which Persson and Tabellini (1996) identify as a core federal function.

The first units of central provision are free at the margin. Because G is non-rival, a unit of ξ delivers country i the full benefit ϑ/G^N while costing it only half a unit of foregone transfer, so the EC's implicit one-half matching rate leaves a strictly positive net gain at $\xi = 0$. Proposition A.3 shows that this gain is exhausted at an interior level rather than pushing ξ to a corner: $\mathcal{W}_i(\xi)$ is strictly concave, rising at $\xi = 0$ and falling without bound as private consumption is squeezed toward zero, with a unique optimum ξ_i^* at which the expected marginal public-good benefit equals half the expected marginal utility of consumption.

Expanded Pareto frontier. I re-solve the full (θ, τ) grid at $\xi = 0.05$ under the reduced surplus (5.2), so that the $\vartheta \log G$ term feeds back into default decisions, debt policy and bond prices.[22]

Public-good provision reshapes the Pareto landscape. In the baseline S's welfare gain never exceeds 0.006 percent of consumption anywhere on the grid; under public-good provision S gains up to 0.34 percent, and 399 of the 441 grid cells are Pareto improving, spanning the whole range of both instruments up to $\theta = 0.5$. That is two orders of magnitude in the level of S's gain and a fivefold expansion of the region. It is the public good doing the work rather than the wider tax grid: at the same $\tau = 0.015$ and the same $\theta = 0$, where each economy delivers S its largest gain at that rate, the public-good economy is worth 0.197 percent to S against 0.004 percent in the baseline. That comparison also locates the gain, because it is taken at $\theta = 0$. What pays S is the tax and what the tax buys, and none of it requires common debt.

The two thresholds of the baseline change places, which is the sharpest statement of what the public good does. In the baseline the tax that pays S sits an order of magnitude above the tax that keeps the EC budget in surplus: at $\theta = 0.1$, S turns positive at $\tau = 0.012$ against $\bar{\tau}(0.1) = 0.0009$.

[21] Were both countries contributing strictly, each would cut its own contribution by exactly the $\xi/2$ it is charged, and the program would be neutral in the sense of Bergstrom et al. (1986). Appendix A.5 states the neutrality result and the conditions under which it fails.

[22] Appendix A.5 reports the closed-form treatment of the local provision game inside the solver, together with a sensitivity exercise that varies its one free constant over a factor of four and leaves every probed cell's welfare sign intact. The level $\xi = 0.05$ sits below the benchmark optimum ξ^* of Proposition A.3, so the gains below are those of a modest common budget rather than of the optimal one.

Figure 5.1: Welfare heatmaps with public-good provision ($\xi = 0.05$)

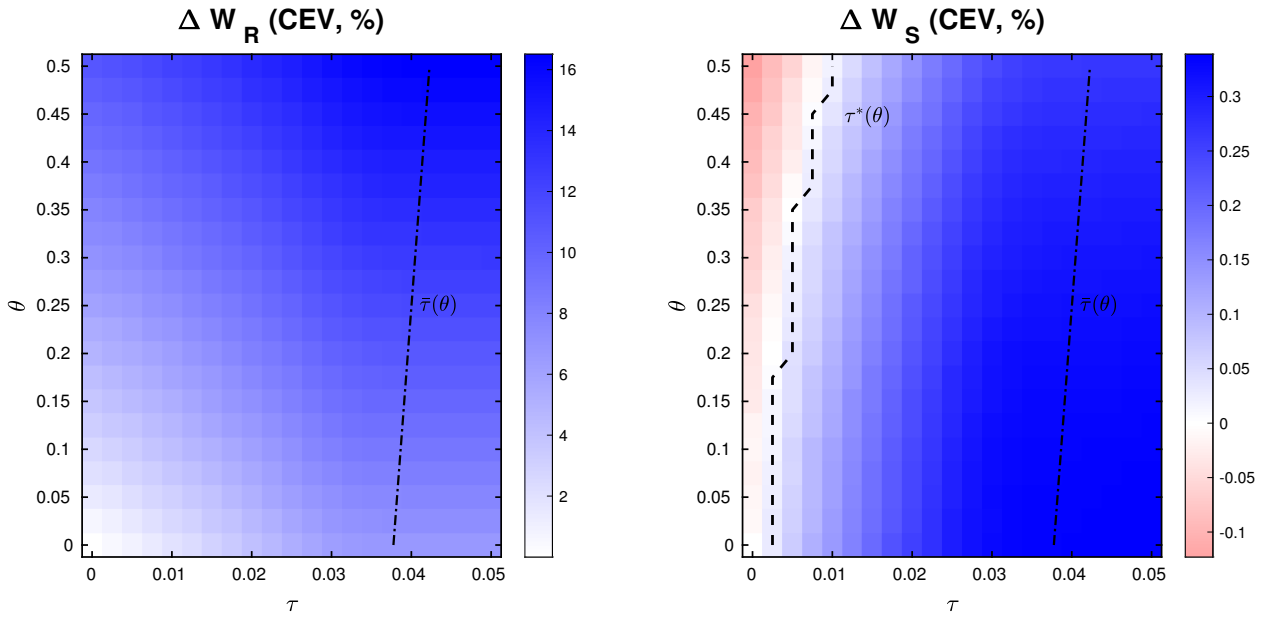

The figure shows the welfare change of each country over the policy grid when the common budget devotes $\xi = 0.05$ to a public good both countries consume, in the same units and format as Figure 4.1. *Left*: $\Delta\mathcal{W}_R(\theta, \tau; \xi = 0.05)$. *Right*: $\Delta\mathcal{W}_S(\theta, \tau; \xi = 0.05)$. Both are measured against the no-integration point $(\theta, \tau) = (0, 0)$ *within* the public-good model, not against the baseline of Section 4. The dash-dotted curve is the public-good threshold $\bar{\tau}^{\xi}(\theta) = \xi / y^{min} + \bar{\tau}(\theta)$ rather than the baseline $\bar{\tau}(\theta)$, and the dashed curve on the right panel is the compensating tax $\tau^*(\theta)$, which here runs well to its left at every θ. The two panels carry independent color scales. The tax grid is extended to $\tau \in [0, 0.05]$, so the horizontal axis is not that of Figure 4.1 and the two should be compared at matched τ.

Under public-good provision the ordering reverses. The threshold itself rises, since ξ has to be funded out of the same revenue, but the compensating tax falls by more: it is 0.0025 for every $\theta \leq 0.175$ and never exceeds 0.010 anywhere on the grid, between a fifteenth and a quarter of $\bar{\tau}^{\xi}(\theta)$.[23]

The safe country is therefore compensated well before the common budget balances. It accepts tax rates at which the EC still runs a deficit in low income states, and so still calls on S to cover part of the shortfall, because the public good bought with the revenue is worth more to S than the share of the shortfall it finances. In the baseline no such trade exists: below $\bar{\tau}$ the revenue buys S nothing it values, so clearing the threshold is a precondition for compensation rather than an alternative to it. That compensation also survives the grant: in the stationary comparison of Appendix A.4, reported for this economy in Figure A.13, S gains in 311 of the 441 cells and by up to 0.34 percent, against at most 0.006 percent in the baseline, while R's gain remains front-loaded and turns negative at nearly every cell.

[23]The boundary is a grid object, uncertain by one tax column at the handful of θ where $\Delta\mathcal{W}_S$ at the first positive cell is itself under 0.006 percent, and two of the 441 points stop at the iteration cap for the reason given in Appendix A.7. Neither qualification touches the reversal, which holds with an order of magnitude to spare.

The two instruments do different jobs, and the public good makes the division explicit. Proposition 3.1 holds in the public-good economy exactly as in the baseline, so θ remains neutral for S above the threshold and every euro that reaches S still arrives through τ. In the baseline that euro has to be raised from R and handed to S, which is why compensation is proportional to τ and why the compensating tax ends up an order of magnitude above the fiscal threshold.

Joint provision breaks that zero-sum structure, and the corner is what makes it create surplus rather than relocate it: the country with the higher marginal utility of consumption contributes nothing of its own and cannot undo the levy. The safe country's participation constraint, which binds throughout the baseline and confines the Pareto region to $\theta \leq 0.125$, is therefore slack almost everywhere once $\xi > 0$.

Common taxation and spending is what makes the arrangement mutually beneficial; common borrowing is what makes the risky country accept it. R's gain runs almost entirely through θ and S's entirely through τ in both economies, and the public good does not join the two legs together. It enlarges the leg that pays S until S's participation constraint stops binding at any scale of the leg that pays R.

The requirement that common debt be paired with non-negligible common taxation is therefore a statement about a facility that only issues and rebates, and it is substantially weakened once the common budget also buys something both members want, which is closer to what the European budget actually does. The corollary is that a common budget which taxes and spends is Pareto improving on its own, at $\theta = 0$: the case for the debt leg rests on the risky country's gain from front-loading, not on anything that leg delivers to the safe country.

5.2 Uncertainty over Permanent Fiscal Roles

A one-in-ten probability of a reversal in fiscal roles more than doubles the Pareto-improving region and nearly doubles the largest program the safe country will accept. The division of the euro area into a permanently risky periphery and a permanently safe core is a modeling simplification, and relaxing it strengthens the ex-ante case for fiscal integration.

I keep the roles permanent and make the assignment itself uncertain. At date 0, before any national state is inherited, one of two configurations is drawn once and holds forever:

- **Normal configuration** (N), drawn with probability $\pi_{\mathsf{N}} = 1 - \pi_{\mathsf{F}}$: country R is risky and S is safe, exactly as in Section 2.
- **Flipped configuration** (F), drawn with probability π_{F}: the roles are reversed, S is risky (impatient, strategic defaulter) and R is safe. The preferences, default cost, and bond-pricing equations of configuration F are the mirror image of configuration N.

The draw allocates a complete permanent economy rather than a label: preferences, default technology, national balance sheet, stationary distribution and pricing schedule travel together, and no

inherited debt stock crosses a change of role.[24] That is a strong reading of a role assignment, and it is what makes the two configurations exact mirror images.

Because each configuration is permanent, the model within it is the baseline of Section 2 with the two countries relabeled, and nothing is left unsolved by the drawing of the role. The mirror mapping is exact when R and S are symmetric in preferences and technology, so that each is capable of assuming either role; the qualification after Proposition 5.1 sets out what changes without it. Writing $\mathcal{W}_i^s(\theta,\tau)$ for the ergodic welfare of country $i \in \{R,S\}$ in configuration $s \in \{\mathsf{N},\mathsf{F}\}$, the symmetric construction gives

$$\mathcal{W}_R^{\mathsf{F}}(\theta,\tau) = \mathcal{W}_S^{\mathsf{N}}(\theta,\tau), \qquad \mathcal{W}_S^{\mathsf{F}}(\theta,\tau) = \mathcal{W}_R^{\mathsf{N}}(\theta,\tau). \tag{5.3}$$

Ex-ante welfare. The ex-ante welfare of country i, before the configuration is drawn, is

$$\mathcal{W}_i^{\mathrm{ea}}(\theta,\tau) \;=\; \pi_{\mathsf{N}}\,\mathcal{W}_i^{\mathsf{N}}(\theta,\tau) \;+\; \pi_{\mathsf{F}}\,\mathcal{W}_i^{\mathsf{F}}(\theta,\tau). \tag{5.4}$$

The relevant welfare differences are the raw utility differences $D_i \equiv \mathcal{W}_i^{\mathsf{N}}(\theta,\tau) - \mathcal{W}_i^{\mathsf{N}}(0,0)$, and I report the entry threshold of (5.5), a ratio of such differences and so free of the consumption unit. Write $\mathcal{P}^{\mathsf{N}}$ for the set of policies that are Pareto improving in configuration N and $\mathcal{P}^{\mathrm{ea}}(\pi_{\mathsf{F}})$ for the ex-ante set.

Proposition 5.1 (Ex-Ante Expansion of the Pareto Set under Role Uncertainty)**.** *Suppose the preferences and default technology of R and S are symmetric across the two configurations, so that* (5.3) *holds, and let the prior probability of the flipped configuration satisfy $\pi_{\mathsf{F}} \in (0,\frac{1}{2})$. Then $\mathcal{P}^{\mathsf{N}} \subseteq \mathcal{P}^{\mathrm{ea}}(\pi_{\mathsf{F}})$, and a policy with $D_S < 0 < D_R$ belongs to $\mathcal{P}^{\mathrm{ea}}(\pi_{\mathsf{F}}) \setminus \mathcal{P}^{\mathsf{N}}$ if and only if*

$$\frac{-D_S}{D_R - D_S} \;<\; \pi_{\mathsf{F}} \;<\; \frac{D_R}{D_R - D_S}. \tag{5.5}$$

For policies with $D_R > -D_S$, entry is monotone over $(0,\frac{1}{2})$, so the ex-ante set is nested increasing on the subset satisfying that condition.

Proof. See Appendix Section A.1. □

The mechanism is that S, the binding side of the constraint throughout the baseline, now holds a claim on R's much larger gain, so the left inequality in (5.5) does the work and the right one

[24] Appendix A.1 sets out three features of this construction that bound what the exercise can say: the two configurations are labeled rather than exchangeable, so π_{F} is a prior weight and not the chance of drawing the risky role in a symmetric lottery; the draw is once and for all, so the exercise is silent on a country that carries its own debt stock into the other role; and averaging across configurations requires a cardinal convention on the two roles, which is why the results below are stated as raw utility differences rather than in consumption-equivalent units.

never binds at these magnitudes. Evaluated on the calibrated surfaces, the region grows steadily (Table 5.1).

Table 5.1: Ex-ante Pareto region under role uncertainty

The table shows how the Pareto-improving region grows once each country attaches a prior probability to drawing the risky role. It reports the number of (θ, τ) cells, out of the 441 of the baseline policy grid, satisfying the ex-ante Pareto condition (5.5), and the largest θ any of them reaches, as a function of the prior probability π_F of the flipped configuration. The column $\pi_F = 0$ is the baseline of Section 4, in which the normal assignment is certain, and $\pi_F = 0.34$ is just above the value 0.332 at which the last policy enters; the count of 440 there is every cell of the grid except the no-program cell $(0, 0)$ itself, at which both gains are zero by construction. Entries re-weight the welfare surfaces of Figure 4.1 according to (5.4) and involve no re-solving of the model.

π_F	0	0.05	0.10	0.20	0.34
Ex-ante Pareto-improving cells	78	113	167	328	440
Largest θ	0.125	0.175	0.225	0.400	0.500

The last policy to enter is the largest untaxed program on the grid, $\tau = 0$ at the top of the θ axis. There S's loss is 33 percent of the total gain the two roles divide between them, $-D_S/(D_R - D_S) = 0.332$, so (5.5) admits it only once π_F reaches that value. Every other policy enters well short of it.

Role uncertainty relaxes the binding participation constraint rather than raising both payoffs. The ex-ante payoff of the country favored by the draw to be *safe* rises, but that of the country favored to be risky falls, since it is now averaged against the smaller flipped-configuration value.[25]

5.3 Convenience Yields and Multiple Equilibria

Convenience yields can drive a gap between interest rates of issuers with the same default probability, and EU common debt currently trades at a non-negligible yield premium over comparably rated core euro-area sovereigns (Pallara et al., 2025; Bonfanti and Marcucci, 2026). The baseline sets $r_J = r_S$, implicitly assuming safe-asset status; this subsection lets r_J fall with program scale, and shows that doing so makes the equilibrium scale of integration indeterminate. Throughout, each θ indexes an alternative *one-shot* launch whose entire issue is priced at the rate its announced final scale commands, so a larger program lowers the service cost on every unit it issues.[26]

I replace the constant EC bond rate r_J with a decreasing function of the per-country transfer θ:

$$r_J(\theta) \;=\; \max\{r_S + \eta - \nu\theta,\ \bar{r}_J\}\,, \qquad \eta > 0,\ \nu > 0,\ \ 0 < \bar{r}_J < r_S, \tag{5.6}$$

[25]The symmetry (5.3) is what makes the two configurations one solved model rather than two. Under asymmetry the flipped configuration has to be solved separately, but the argument survives verbatim with D_S and D_R in (5.5) replaced by each country's own flipped-configuration gain: what the proof uses is only that S holds a claim on a gain larger than its own loss, not that the claim is R's.

[26]The incremental protocol below is therefore bargaining over progressively larger pre-launch proposals, not the addition of tranches to an outstanding stock.

where η is the initial illiquidity premium, ν the market-depth elasticity, and $\bar{r}_J$ a long-run floor. The condition $\bar{r}_J < r_S$ ensures that a large enough program earns a convenience yield, while $\bar{r}_J > 0$ rules out above-par pricing. For $\theta \geq \theta^\dagger \equiv (r_S + \eta - \bar{r}_J)/\nu$, the floor binds. The exercise is theoretical: (5.6) is not estimated and the calibrated model is not re-solved under it, so the results characterize shapes rather than measure them, and the figures use illustrative parameters.

Making $\bar{q}$ a function of θ breaks the cash-flow neutrality that leaves S indifferent to the scale of the program in the baseline, and it breaks it in both directions, because S discounts the EC's service stream at its own rate q_S while the EC funds itself at $\bar{q}(\theta)$. The two coincide only at $\theta = \eta/\nu$: below that scale the illiquidity premium makes the program cost S more than the cash it delivers, above it the convenience yield reverses the comparison.[27]

S's welfare inherits the shape. The fee factor $(1-\bar{q}(\theta))\theta$, of which the per-unit cost is $1-\bar{q}(\theta)$, is humped for the same reason the threshold is: a small program pays the illiquidity premium on every unit it issues, and only a large one earns the convenience yield that pays for it. Proposition A.5 turns this into a *valley*, with $d\mathcal{W}_S/d\theta$ negative near zero, positive near the floor $\theta^\dagger$, and zero at an interior trough θ^v, provided the initial premium η is large enough to dominate the grant at $\theta = 0^+$ and the residual fee at the floor is small enough not to. S's welfare therefore dips below the benchmark at intermediate scales before turning back up (Figure 5.2), so incremental expansion passes through a region S would veto step by step even though a discrete jump past the recovery threshold θ^{rec} can be a Pareto improvement.

Multiple integration equilibria. The valley does more than complicate incremental reform: it renders the equilibrium *scale* of integration indeterminate. Let $\theta^{\max}$ be an exogenous upper bound on the scale of the program, hold the tax fixed at a rate τ that clears the threshold over that whole range, $\tau \geq \max_{\theta\in[0,\theta^{\max}]} \bar{\tau}(\theta)$, as Proposition A.5 requires, and write $\Delta\mathcal{W}_i(\theta) \equiv \mathcal{W}_i(\theta,\tau) - \mathcal{W}_i(0,\tau)$, the gain from common debt at an unchanged tax.[28] Country R gains monotonically, while $\Delta\mathcal{W}_S$ traces the valley and turns non-negative only past the *recovery threshold* $\theta^{\text{rec}} > \theta^v$ at which it crosses back through zero: the trough is where S's welfare stops falling, the recovery threshold where it regains the tax-only benchmark. Consider a decentralized protocol in which the scale is raised in small increments, each requiring both countries to remain weakly better off than the tax-only union,

[27] The fiscal threshold inherits the same shape, becoming non-monotone in θ rather than linear, hump-shaped before the floor binds and linearly increasing after and crossing the baseline threshold at that same scale η/ν; Proposition A.4 in Appendix A.6 states it. Nothing below uses it, since the tax is held fixed above $\bar{\tau}$ over the whole range.

[28] The comparator is $(0,\tau)$ rather than autarky $(0,0)$ because that is what the protocol's withdrawal option delivers: abandoning common debt at a fixed tax leaves the symmetric tax and rebate in place. It also makes $\Delta\mathcal{W}_i(0) = 0$ hold by construction, which a comparison against $(0,0)$ would not, since Section 4.2 finds a tax-only strip at $\theta = 0$ on which S already gains. Fixing τ rather than tracking $\bar{\tau}(\theta)$ also makes $d\mathcal{W}_S/d\theta$ a total derivative along the path, so Proposition A.5 applies to it unchanged; along the threshold the derivative would carry the additional $\bar{\tau}'(\theta)\,\partial\mathcal{W}_S/\partial\tau$ terms of (3.15), whose sign flips at the peak θ^{**}. Nothing in the multiplicity argument turns on which path is used.

Figure 5.2: Welfare changes under variable convenience yield ($\tau \geq \bar{\tau}(\theta)$)

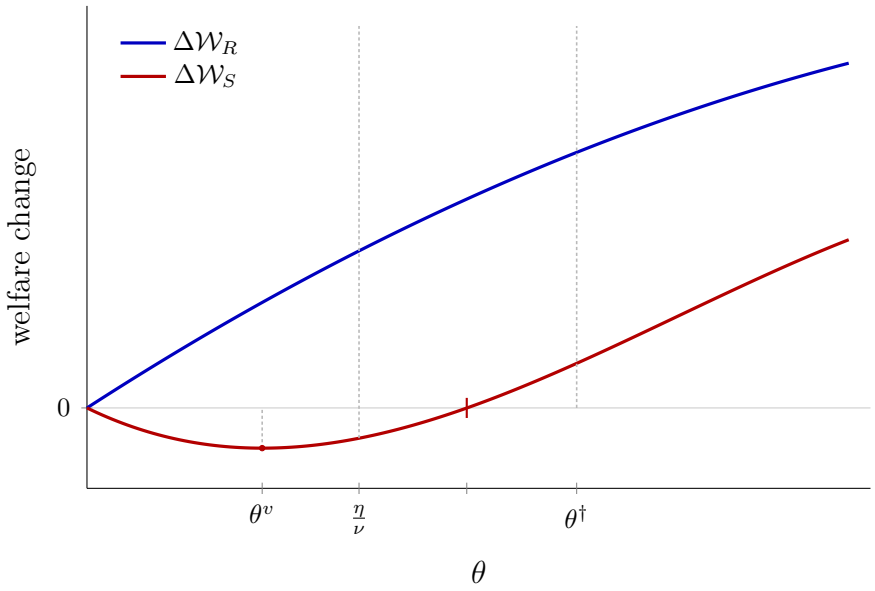

The figure shows the welfare change of the two countries as a function of the scale of the program when the price at which the common institution issues varies with that scale, as in (5.6). Both curves are measured relative to no common debt at the same tax ($\theta = 0$, τ unchanged), at a tax rate held fixed above $\bar{\tau}(\theta)$ over the whole range. $\Delta\mathcal{W}_R$ (blue) is monotonically increasing. $\Delta\mathcal{W}_S$ (red) exhibits the welfare valley of Proposition A.5, dipping below zero for small θ and recovering beyond the trough θ^v; the tick mark on the horizontal axis is its zero-crossing. The trough occurs before η/ν, where $r_J = r_S$, consistent with Remark A.2: the recovery is initiated by the saving the scale economy delivers on inframarginal units, while EC bonds are still more expensive than national ones, and the spillover benefit only reinforces it. The horizontal landmarks η/ν and $\theta^\dagger$ are placed at the values implied by the parameters of Figure A.2, but the two curves are illustrative shapes drawn by hand rather than functions computed from those parameters or from the calibrated model, so only the shapes and the ordering of θ^v, η/ν and $\theta^\dagger$ are meaningful.

with either being free to withdraw to $\theta = 0$ at the same tax.

Proposition 5.2 (Multiple Integration Equilibria)**.** *Suppose that the conditions of Proposition A.5 hold, that $\Delta\mathcal{W}_R$ is strictly increasing on $[0, \theta^{\max}]$, and that the recovery is strong enough that $\Delta\mathcal{W}_S(\theta) > 0$ for some $\theta \leq \theta^{\max}$, so that a recovery threshold $\theta^{\text{rec}} \in (\theta^v, \theta^{\max})$ exists. Then:*

(i) The safe country's participation set $\mathcal{F} \equiv \{\theta \in [0, \theta^{\max}] : \Delta\mathcal{W}_S(\theta) \geq 0\}$ is disconnected, $\mathcal{F} = \{0\} \cup [\theta^{\text{rec}}, \theta^{\max}]$, its two components separated by the welfare valley $(0, \theta^{\text{rec}})$ on which S strictly loses.

(ii) Under the incremental participation-respecting protocol, both no integration $\theta = 0$ and the ambitious program $\theta^{high} = \theta^{\max}$ are stable equilibria.

(iii) θ^{high} weakly Pareto-dominates $\theta = 0$, strictly for R, yet is unreachable from it by any continuous participation-feasible path: the transition requires a discrete jump of size at least θ^{rec}.

Proof. See Appendix Section A.6. □

The two equilibria read as market depth, which connects the model to the safe-asset coordination literature (He et al., 2019). In the low equilibrium the program stays small, $r_J(\theta)$ sits near $r_S+\eta$, and the illiquidity premium keeps common bonds more expensive than national debt, which is what keeps the program small; in the high equilibrium $r_J(\theta)$ reaches the floor $\bar{r}_J$ and the realized convenience yield is what can make the large program mutually beneficial. Expectations about the eventual depth of the common-bond market are thus self-fulfilling through $r_J(\theta)$, and the convenience yield is the source of the multiplicity: at $\eta = 0$ the valley collapses, $\mathcal{F}$ is the whole interval, and the integration gains are monotone with a unique stable scale.

Farhi and Maggiori (2018) obtain multiplicity in a related way, from an issuer whose borrowing cost depends on the size of its own safe-asset supply, but their indeterminacy binds when the stock is large and is a confidence problem among investors. The one here binds when the program is small and is a coordination problem between the two member states, who must agree on the scale before the market that would justify it exists.[29]

The valley thus provides a rationale for *big-bang* fiscal integration, a program of sufficient scale launched at once, consistent with the EU's NextGenerationEU experience, and it explains why moderate proposals face resistance from safe countries. The path to the Pareto region passes through a zone that is strictly worse for S, and only a program large enough to clear both that zone and the elevated threshold phase reaches the other side.

6 Conclusion

Fiscal integration in a currency union need not be a one-way transfer. A common institution that issues common bonds funded by a non-defaultable proportional income tax can raise the welfare of both the fiscally weak and the fiscally strong country, with no compensating transfer outside the common institution. What makes this possible is enforcement and insurance rather than a one-way subsidy: a tax members cannot repudiate is collectible in every state, the institution borrows at the safe rate and the insurance the tax carries keeps paying when default occurs. The weak country's gain is from intertemporal trade and the strong country's from insurance, and because grant and tax base are symmetric, neither is financed by an expected transfer from the other.

The analysis also establishes what the design must deliver. A fiscal threshold separates the tax rates under which the program worsens default incentives from those under which it eliminates them. The decompositions make the mutual benefit transparent: the risky country gains through the initial grant, not cheaper credit, and the safe country only through the common tax, which

[29]The multiplicity lives in the coordination over the program scale θ, a policy object jointly selected by the two governments, not in the recursive equilibrium for a given (θ, τ). Holding the policy fixed, R's default problem and the associated pricing fixed point admit a unique solution, so every point of the grid is itself unambiguous.

channels resources toward whichever country is temporarily poorer. The threshold is small, the compensating tax an order of magnitude larger, but the institutional distance from the status quo smaller than the political debate suggests. The hard part, once the collection power exists, is not enforcing the common claim; it is paying the safe country enough to accept the program.

The two legs of the arrangement do different jobs, and spending part of the common surplus on public goods makes the division plain. Common taxation and spending is what makes the arrangement mutually beneficial, because a jointly consumed good pays both countries at once rather than moving resources between them; common borrowing is what makes the risky country accept it, because a front-loaded grant is worth more to an impatient borrower than the perpetual claim that funds it. The safe country is never paid by the size of the program, only by the tax, and once the budget also buys a non-rival good it is paid at rates below the threshold at which the institution stops running deficits. Paying the strong country then ceases to be the expensive part of the design. The implication for future European programs is that the common budget rather than the common bond is what has to carry the safe country, and that a facility which only issues and rebates asks the tax to do work it can barely do.

The logic has a precedent, whose parallel to the euro area Sargent (2012) develops in detail. Hamilton's 1790 proposal to have the federal government assume the states' war debts was contested on exactly the grounds familiar today, that disciplined states were subsidizing those that had not serviced their obligations. It worked because it was paired with a new federal revenue instrument that made the program self-sustaining, and the strong states gained lower borrowing costs, a stable confederation, and a federal capacity for common goods. It was accepted not as charity but as a credible mutual commitment, and demonstrating the same for Europe is what this paper attempts.

References

Aguiar, M. and Amador, M. (2019). A contraction for sovereign debt models. *Journal of Economic Theory*, 183:842–875.

Aguiar, M., Amador, M., Farhi, E., and Gopinath, G. (2015). Coordination and crisis in monetary unions. *The Quarterly Journal of Economics*, 130(4):1727–1779.

Aguiar, M. and Gopinath, G. (2006). Defaultable debt, interest rates and the current account. *Journal of International Economics*, 69(1):64–83.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3):659–684.

Ando, S., Dell'Ariccia, G., Gourinchas, P.-O., Lorenzoni, G., Peralta-Alva, A., and Roch, F. (2023). Debt Mutualization in the Euro Area: A Quantitative Exploration. IMF Working Paper 2023/059, International Monetary Fund.

Arellano, C. (2008). Default risk and income fluctuations in emerging economies. *American Economic Review*, 98(3):690–712.

Asdrubali, P., Sørensen, B. E., and Yosha, O. (1996). Channels of interstate risk sharing: United states 1963–1990. *The Quarterly Journal of Economics*, 111(4):1081–1110.

Auclert, A. and Rognlie, M. (2016). Unique equilibrium in the Eaton–Gersovitz model of sovereign debt. *Journal of Monetary Economics*, 84:134–146.

Azzimonti, M. and Quadrini, V. (2024). International spillovers and bailouts. *The Review of Economic Studies*, 91(1):77–128.

Bergstrom, T., Blume, L., and Varian, H. (1986). On the private provision of public goods. *Journal of Public Economics*, 29(1):25–49.

Berriel, R., Gonzalez-Aguado, E., Kehoe, P. J., and Pastorino, E. (2024). Is a fiscal union optimal for a monetary union? *Journal of Monetary Economics*, 141:157–177.

Besley, T. and Coate, S. (2003). Centralized versus decentralized provision of local public goods: a political economy approach. *Journal of Public Economics*, 87(12):2611–2637.

Bianchi, F., Fang, Q., Melosi, L., and Rogantini Picco, A. (2026). Who is afraid of eurobonds? Working Paper 35510, National Bureau of Economic Research.

Bonfanti, G. and Marcucci, J. (2026). A European Safe Asset? Not Without the Investors. Questioni di Economia e Finanza (Occasional Papers) 1010, Bank of Italy.

Brunnermeier, M. K., Langfield, S., Pagano, M., Reis, R., Van Nieuwerburgh, S., and Vayanos, D. (2017). ESBies: Safety in the tranches. *Economic Policy*, 32(90):175–219.

Chatterjee, S. and Eyigungor, B. (2012). Maturity, indebtedness, and default risk. *American Economic Review*, 102(6):2674–2699.

Cioffi, M., Rizza, P., Romanelli, M., and Tommasino, P. (2019). Outline of a Redistribution-Free Debt Redemption Fund for the Euro Area. Occasional Paper 479, Bank of Italy.

Cruces, J. J. and Trebesch, C. (2013). Sovereign defaults: The price of haircuts. *American Economic Journal: Macroeconomics*, 5(3):85–117.

Dávila, E. and Schaab, A. (2025). Welfare Assessments with Heterogeneous Individuals. *Journal of Political Economy*, 133(9):2918–2961.

Dávila, E. and Weymuller, C.-H. (2017). Optimal Joint Bond Design. Working paper.

de Ferra, S. and Mallucci, E. (2025). Avoiding sovereign default contagion: A normative analysis. *Journal of International Economics*, 154:104040.

Draghi, M. (2024). The Future of European Competitiveness. Part A: A Competitiveness Strategy for Europe. Technical report, European Commission, Brussels.

Eaton, J. and Gersovitz, M. (1981). Debt with potential repudiation: Theoretical and empirical analysis. *The Review of Economic Studies*, 48(2):289–309.

Farhi, E. and Maggiori, M. (2018). A model of the international monetary system. *The Quarterly Journal of Economics*, 133(1):295–355.

Farhi, E. and Werning, I. (2017). Fiscal unions. *American Economic Review*, 107(12):3788–3834.

Ferrari, A. and Rogantini Picco, A. (2023). Risk sharing and the adoption of the euro. *Journal of International Economics*, 141:103727.

Hatchondo, J. C., Martinez, L., and Önder, Y. K. (2017). Non-defaultable debt and sovereign risk. *Journal of International Economics*, 105:217–229.

He, Z., Krishnamurthy, A., and Milbradt, K. (2019). A model of safe asset determination. *American Economic Review*, 109(4):1230–62.

Jarociński, M. and Maćkowiak, B. (2018). Monetary-fiscal interactions and the euro area's malaise. *Journal of International Economics*, 112:251–266.

Karaivanov, A., Mojon, B., da Silva, L. P., Tejada, A. P., and Townsend, R. M. (2026). Incentive-compatible unemployment reinsurance for the euro area. *European Economic Review*, 186:105325.

Leandro, Á. and Zettelmeyer, J. (2019). Creating a Euro area safe asset without mutualizing risk (much). *Capital Markets Law Journal*, 14(4):488–517.

Lucas, R. E. (1987). *Models of Business Cycles*. Basil Blackwell, Oxford.

Martinez, J., Philippon, T., and Sihvonen, M. (2022). Does a currency union need a capital market union? *Journal of International Economics*, 139:103675.

Oates, W. E. (1999). An essay on fiscal federalism. *Journal of Economic Literature*, 37(3):1120–1149.

Pallara, K., Pericoli, M., and Tommasino, P. (2025). Issuing European Safe Assets: How to Get the Most Out of Eurobonds? Occasional Paper 937, Bank of Italy.

Pallara, K. and Renne, J.-P. (2024). Fiscal limits and the pricing of eurobonds. *Management Science*, 70(2):1216–1237.

Persson, T. and Tabellini, G. (1996). Federal fiscal constitutions: Risk sharing and redistribution. *Journal of Political Economy*, 104(5):979–1009.

Samuelson, P. A. (1954). The pure theory of public expenditure. *The Review of Economics and Statistics*, 36(4):387–389.

Sargent, T. J. (2012). Nobel Lecture: United States Then, Europe Now. *Journal of Political Economy*, 120(1):1–40.

Schmitt-Grohé, S. and Uribe, M. (2003). Closing small open economy models. *Journal of International Economics*, 61(1):163–185.

Tirole, J. (2015). Country solidarity in sovereign crises. *American Economic Review*, 105(8):2333–2363.

Yue, V. Z. (2010). Sovereign default and debt renegotiation. *Journal of International Economics*, 80(2):176–187.

Appendix A Additional model results

A.1 Model proofs

Proof of Proposition 2.1. (i) For any $\tau \geq \bar{\tau}(\theta) = 8(1-\bar{q})\theta/y^{min}$,

$$\begin{aligned}\omega_t &= \tau(y_{R,t} + y_{S,t}) - 8(1-\bar{q})\theta \\ &\geq \bar{\tau}(\theta)\,(y_{R,t} + y_{S,t}) - 8(1-\bar{q})\theta \\ &\geq \bar{\tau}(\theta)\,y^{min} - 8(1-\bar{q})\theta \;=\; 0,\end{aligned}$$

where the second line uses $\tau \geq \bar{\tau}(\theta)$ and the third uses $y_{R,t} + y_{S,t} \geq y^{min}$ together with $\bar{\tau}(\theta) \geq 0$. Since $\omega_t \geq 0$, the additional fiscal cost on S in default periods is $\max\{-\frac{1}{2}\omega_t, 0\} = 0$, so S's total exposure to R's default is exactly φ. Under repayment, R's EC receipt is $\frac{1}{2}\omega_t \geq 0$; under default, R's EC receipt is $\max\{\frac{1}{2}\omega_t, 0\} = \frac{1}{2}\omega_t \geq 0$. The two are identical, so the fiscal channel contributes zero to the default incentive.

(ii) If $\tau \geq \bar{\tau}(\theta + \Delta\theta)$, then $\omega_t \geq 0$ holds under the expanded program by (i), and S's fiscal exposure remains exactly φ. □

Proof of Proposition 3.1. Under $\beta_S(1+r_S) = 1$ and a slack borrowing constraint, S's Euler equation for the risk-free bond holds with equality in every period:

$$u'_S(c_{St}) \;=\; \beta_S(1+r_S)\,\mathbb{E}_t[u'_S(c_{S,t+1})] \;=\; \mathbb{E}_t[u'_S(c_{S,t+1})].$$

Marginal utility is therefore a martingale under the household's beliefs. Iterating the tower property from $t=0$ gives $\mathbb{E}_0[u'_S(c_{St})] = u'_S(c_{S0})$ for every $t \geq 0$. Dividing by $u'_S(c_{S0})$ and multiplying by β_S^t:

$$\mathbb{E}_0[\Pi_{St}] \;=\; \beta_S^t \;=\; \bar{q}^t, \qquad t \geq 0,$$

where the second equality uses $\beta_S = q_S = \bar{q}$. Summing the EC fee flow $4(1-\bar{q})$ per unit of θ from $t=0$, the first period in which the EC's budget operates, and using the geometric series:

$$4(1-\bar{q})\sum_{t=0}^{\infty}\mathbb{E}_0[\Pi_{St}] \;=\; 4(1-\bar{q})\sum_{t=0}^{\infty}\bar{q}^{\,t} \;=\; 4(1-\bar{q})\cdot\frac{1}{1-\bar{q}} \;=\; 4,$$

which is (3.11). □

The proof assumes a slack band, as Proposition 3.1 states, and the two barriers cut in opposite directions when that fails. If the borrowing limit $b'_S \leq \bar{b}_S$ binds with positive probability, the Euler equation holds as $u'_S(c_{St}) \geq \mathbb{E}_t[u'_S(c_{S,t+1})]$, marginal utility is a strict supermartingale, the chain of equalities above becomes $\sum_{t=0}^{\infty}\mathbb{E}_0[\Pi_{St}] \leq 1/(1-\bar{q})$, and the present value of the fees falls weakly

below the grant, a residual favorable to S. If instead the deleveraging limit $b'_S \geq \underline{b}_S$ binds, S cannot carry the whole grant forward as a permanently lower liability, marginal utility is a submartingale, and the inequalities reverse: the present value of the fees exceeds the grant. It is this second case that operates in the calibration, and footnote 16 measures it.

Proof of Proposition 5.1. Apply (5.3) to (5.4). The ex-ante gains of the two countries are

$$\mathcal{W}_R^{\text{ea}}(\theta,\tau) - \mathcal{W}_R^{\text{ea}}(0,0) = \pi_{\mathsf{N}} D_R + \pi_{\mathsf{F}} D_S, \qquad \mathcal{W}_S^{\text{ea}}(\theta,\tau) - \mathcal{W}_S^{\text{ea}}(0,0) = \pi_{\mathsf{N}} D_S + \pi_{\mathsf{F}} D_R,$$

since the flipped-configuration value of each country is the normal-configuration value of the other.

Containment. If $(\theta,\tau) \in \mathcal{P}^{\mathsf{N}}$ then $D_R > 0$ and $D_S > 0$, so both expressions are convex combinations of strictly positive numbers and $(\theta,\tau) \in \mathcal{P}^{\text{ea}}(\pi_{\mathsf{F}})$ for every $\pi_{\mathsf{F}} \in [0,1]$.

Strict entry. Let $D_S < 0 < D_R$, so $(\theta,\tau) \notin \mathcal{P}^{\mathsf{N}}$ and $D_R - D_S > 0$. Writing $\pi_{\mathsf{N}} = 1 - \pi_{\mathsf{F}}$ and collecting terms, $\pi_{\mathsf{N}} D_S + \pi_{\mathsf{F}} D_R > 0$ is equivalent to $\pi_{\mathsf{F}}(D_R - D_S) > -D_S$, and $\pi_{\mathsf{N}} D_R + \pi_{\mathsf{F}} D_S > 0$ to $\pi_{\mathsf{F}}(D_R - D_S) < D_R$. Together these are exactly (5.5). Since the left bound is strictly positive and the right bound exceeds $\frac{1}{2}$ whenever $D_R > -D_S$, raising π_{F} over $(0, \frac{1}{2})$ admits policies and never removes them, so $\mathcal{P}^{\text{ea}}(\pi_{\mathsf{F}})$ is nested increasing there. □

Scope of the role-uncertainty exercise. Three features of the construction bound what Proposition 5.1 can say. First, the two configurations are labeled rather than exchangeable, so π_{F} is a prior weight on the mirror allocation and not the chance that either particular country draws the risky role in a symmetric lottery; the latter would force $\pi_{\mathsf{F}} = 1/2$. Second, the draw is once and for all. It is not a hazard of switching roles later, and the exercise is silent on the transition that would make such a switch interesting: a country carrying its own inherited debt stock into the other role, with lenders pricing that possibility into both countries' bonds today. That is a different model rather than a reweighting of this one, and it is left for future work. It matters because the two roles carry very different balance sheets at the calibrated origin, 104.5 percent of annual output on the risky side against 61.5 percent on the safe side, so a country that genuinely changed role would arrive with the wrong one of the two.

Third, averaging across configurations requires a common cardinal scale for the two roles, since the prizes of the lottery assign the recipient different preferences. I use the same CRRA felicity function in the same consumption units for both roles, with only the discount factor role-specific, exactly as calibrated. This is a normative convention about the ex-ante criterion rather than an implication of the two solved national problems. Section 5.2 accordingly works with the raw utility differences D_i and reports the entry threshold of (5.5), a ratio of such differences and so free of the consumption unit, rather than restating the exercise in the consumption-equivalent units of Section 3.1. A consumption equivalent of (5.4) is available and is exact, since scaling consumption by $1+\kappa$ in both configurations multiplies each $\mathcal{W}_i^s$ by $(1+\kappa)^{1-\gamma}$ and hence $\mathcal{W}_i^{\text{ea}}$ by the same factor.

It is not, however, comparable with the consumption equivalents reported elsewhere in the paper: the mixture adds two lifetime utilities generated under discount factors an order of magnitude apart in implied horizon, so the resulting figure is dominated by that difference in scale rather than by the size of the welfare gain. Concretely, the ex-ante consumption equivalent of R at $(\theta, \tau) = (0.5, 0)$ falls from 9.80 percent at $\pi_{\mathsf{F}} = 0$ to 2.63 percent at $\pi_{\mathsf{F}} = 0.05$ and 0.29 percent at $\pi_{\mathsf{F}} = 0.34$, not because R's gain shrinks but because the denominator inherits the safe role's much larger lifetime utility level. The signs are unaffected, which is all the participation condition uses.

A.2 Pareto-improvement conditions

Section 3 establishes that S is the binding side of the Pareto constraint along both instruments. This subsection states and proves the two conditions that follow. They run in opposite directions: along θ the available result is *sufficient* and rests on an assumption about how R's default rule responds to the program, while along τ it is *necessary* and rests on the sign of the risk-sharing terms.

Proposition A.1 (Local Sufficient Conditions for a Non-Negative θ-Attribution to Safe-Country Welfare)**.** *Fix an interval of program sizes $[\theta_0, \theta_1]$ and a differentiable policy path $\tau = \tau(\theta)$ defined on it, and let*

$$\mathcal{A}_S(\theta) \;\equiv\; \frac{1}{u'_S(c_{S0})}\left[\frac{\partial \mathcal{W}_S}{\partial \theta} + \tau'(\theta)\,\frac{\partial \mathcal{W}_S}{\partial \tau}\right] \tag{A.1}$$

denote the marginal welfare attribution to S along that path, the two partial derivatives being those of (3.10) *and* (3.15) *and inheriting their convention. Assume that for every $\theta \in [\theta_0, \theta_1]$:*

(a) $\tau(\theta) \geq \bar{\tau}(\theta)$ (so Proposition 2.1(i) applies and S's fiscal exposure is exactly φ).

(b) The EC surplus is weakly non-decreasing in θ state-by-state along the path, i.e. $d\omega_t/d\theta \geq 0$:

$$(y_{R,t} + y_{S,t})\,\tau'(\theta) - 8(1 - \bar{q}) \geq 0 \quad \forall (y_{R,t}, y_{S,t}) \text{ on the support.} \tag{A.2}$$

A sufficient condition is $\tau'(\theta) \geq 8(1 - \bar{q})/y^{min}$, which also keeps (a) satisfied as θ grows.

(c) S's SDF-weighted risk-sharing flow is non-negative along the path:

$$E^{all}\left[\Pi_{St}\,\frac{y_{Rt} - y_{St}}{2}\right] \;\geq\; 0. \tag{A.3}$$

(d) R's default rule is weakly decreasing in θ along the path, state by state: for every state $(b, y) = (b, y_R, y_S)$ on the ergodic support and every $\theta_a < \theta_b$ in $[\theta_0, \theta_1]$,

$$d_R(b, y;\, \theta_b, \tau(\theta_b)) \;\leq\; d_R(b, y;\, \theta_a, \tau(\theta_a))\,. \tag{A.4}$$

The condition is imposed along the path*, so it restricts the joint response of the default rule to the two instruments together rather than the response to θ at fixed τ. When the default rule takes a threshold form in R's own income, $d_R(b, y; \theta, \tau) = \mathbf{1}\{y_R < y_R^*(b, y_S; \theta, \tau)\}$,* (A.4) *is the requirement that $y_R^*(b, y_S; \theta, \tau(\theta))$ be weakly decreasing in θ for every (b, y_S). The threshold has to be indexed by y_S as well as by b: the EC surplus* (2.4) *places aggregate income in both of R's budget constraints, so S's income is an argument of R's repayment and default values whenever $\tau > 0$, which is the range this proposition covers.*

Then $\mathcal{A}_S(\theta) \geq 0$ at every $\theta \in [\theta_0, \theta_1]$.

Two things are deliberately not claimed. The conclusion is about the decomposition of footnote 9, which holds the invariant distribution fixed, and not *about the exact total derivative of $\mathcal{W}_S$: signing that would additionally require perturbing μ^*, and Section 4.3 measures what the omission is worth once the attribution is integrated over a finite reform. The conclusion is also confined to this interval and to paths that stay above the fiscal threshold throughout it, and carries no implication for program sizes at which (d) fails. Within those limits (d) does more than it is asked for: it signs R's default rule state by state, and the proof uses that statewise form directly rather than the income-conditional rate of Section 3.2.*

Proof. By Proposition 2.1(i), condition (a) implies $\omega_t \geq 0$ in all states along the path, so $\mathcal{D}^- = \emptyset$ and S's EC-deficit-sharing term is identically zero; condition (b) preserves that property as the program grows, since a surplus that is non-negative at θ_0 and weakly rising cannot turn negative.

Condition (d) is statewise monotonicity of R's default decision along the path, so with the invariant distribution held fixed as in footnote 9 it signs the induced change

$$\delta_t^{\text{path}}(s_t) \;\equiv\; -\frac{d\, d_R(s_t)}{d\theta}\Big|_{\tau=\tau(\theta)} \;\geq\; 0 \tag{A.5}$$

at every state $s_t = (b_t, y_t)$ on the ergodic support, read as a local comparative static on a discrete rule. Three features of this step matter. First, (d) signs the *sum* of the two instruments' effects: neither component is signed on its own, which is why the condition has to be stated along the path. Second, the sign is obtained at the level of the state, which is where the spillover channel of (3.10) and (3.15) is formed before anything is averaged. The proof therefore never passes through the income-conditional rate Δ_t and does not inherit the conditional-averaging approximation of footnote 10: because $\varphi \geq 0$ and $\Pi_{St} > 0$, the exact state-level integrand $\Pi_{St}\, \varphi\, \delta_t^{\text{path}}$ is non-negative pointwise, so its expectation is non-negative whatever the joint distribution of the two debt stocks. Third, what is signed is the response of the rule at a fixed distribution. Statewise monotonicity does not on its own sign the unconditional default probability of the reformed economy, because that probability also integrates the rule against an invariant distribution that the program itself relocates, and the two movements need not agree in sign.

The remaining terms must be collected in expectation rather than state by state, because S's net EC position, $-\tau(\theta)y_{St} + \frac{1}{2}\omega_t$, moves with θ by $\frac{1}{2}\tau'(\theta)(y_{Rt} - y_{St}) - 4(1-\bar{q})$ per unit, which is negative in states where S draws the higher income however large τ' is. Condition (b) signs the surplus ω_t, not that net position. In the first term of $\mathcal{A}_S$ the date-0 grant and the present value of the EC fee cancel exactly by Proposition 3.1, leaving the spillover channel alone; the second contributes the risk-sharing term *and* a second spillover term, since τ moves R's default frequency as well. Collecting the two spillover terms into the single path response (A.5),

$$\mathcal{A}_S(\theta) \;=\; \varphi\, E^{\text{all}}\Big[\Pi_{St}\,\delta_t^{\text{path}}\Big] \;+\; \tau'(\theta)\, E^{\text{all}}\Big[\Pi_{St}\,\frac{y_{Rt}-y_{St}}{2}\Big] \;\geq\; 0,$$

the first term because $\delta_t^{\text{path}} \geq 0$ pointwise by (d) and $\varphi \geq 0$, the second because $\tau'(\theta) \geq 0$ by (b) and the expectation is non-negative by (c). Since θ was an arbitrary point of $[\theta_0, \theta_1]$, $\mathcal{A}_S \geq 0$ throughout the interval. Whether $\mathcal{W}_S$ itself is weakly increasing there is the same question one step further out, and it turns on the displacement of μ^* that $\mathcal{A}_S$ omits by construction. □

Conditions (a) and (b) restrict the policy path and can be imposed by construction. Conditions (c) and (d) cannot, since both are properties of the equilibrium rather than of the policy. The calibrated economy satisfies (c) comfortably, S's risk-sharing term being the largest of its entries in Table 4.2 and the only one of any size, but it rejects (d), both at $\tau = 0$ and at $(\theta, \tau) = (0.1,\, 0.01)$, a cell inside the Pareto region and an order of magnitude above $\bar{\tau}$ (Section 4.3). The rejection also shows how far the two movements separated in the proof can pull apart. Raising θ from 0 to 0.5 at $\tau = 0$ lowers R's ergodic debt ratio from 104.5 to 99.8 percent of annual output, a displacement of the invariant distribution that on its own works against default, and yet the annualized default frequency rises from 1.09 to 1.15 percent: the default rule moves adversely by more than the distribution moves favorably. The failure of (d) is therefore understated by the aggregate frequency, not manufactured by it. The proposition isolates the condition under which θ on its own would carry S, together with the finding that this economy does not meet it. Note also that (c) is what does the work when τ rises with θ: the increment to S's welfare travels through the tax, not through the size of the program.

Proposition A.2 (Necessary Conditions for a Non-Negative τ-Attribution)**.** *Suppose $\tau > \bar{\tau}(\theta)$, so that* (3.13)–(3.15) *apply and all behavioral effects on S operate only through φ. A marginal increase in τ raises the welfare attribution of both countries only if both of the following hold:*

$$E^{all}\Big[\Pi_{Rt}\,\frac{y_{St}-y_{Rt}}{2}\Big] + E^{\mathcal{N}}\Big[\Pi_{Rt}\,b'_{Rt}\,\frac{\partial q_R}{\partial \tau}\Big] \;\geq\; 0, \tag{A.6}$$

$$E^{all}\Big[\Pi_{St}\,\frac{y_{Rt}-y_{St}}{2}\Big] + \varphi\, E^{all}[\Pi_{St}\,\Delta_t^{\tau}] \;\geq\; 0, \tag{A.7}$$

with $\Delta_t^\tau = -\partial \Pr_t(d_R = 1)/\partial\tau$ *as in* (3.15). *The first condition requires the SDF-weighted risk-sharing term, augmented by the credit-response channel, to be non-negative for* R. *The risk-sharing component* $E^{all}[\Pi_{Rt}(y_{St} - y_{Rt})/2] \geq 0$ *holds when* R*'s income tends to be low in states where* R*'s marginal utility is high; the credit-response term is non-negative whenever* $\partial q_R/\partial\tau \geq 0$, *i.e. when a higher* τ *disciplines default and raises* R*'s bond price in repayment states. The second condition requires* S*'s risk-sharing term, whose sign is calibration-dependent, to cover any loss on the spillover channel through* φ.

The same scope restriction applies here as in Proposition A.1, and for the same reason. What the two conditions are necessary for is non-negativity of the attributions of footnote 9, which hold the invariant distribution fixed; they are necessary for an actual Pareto improvement only to the extent that those attributions track the exact total derivatives, and Section 4.3 measures what the qualification is worth once an attribution is integrated over a finite reform. Nothing here bounds the set of tax rates satisfying the two conditions: whether it is an interval, and how far it extends, is a quantitative question addressed on the tax grid of Section 4.2.

Proof. Take (3.13)–(3.15), whose left-hand sides are the attributions of footnote 9 rather than exact total derivatives. For the attribution of (3.13) to be non-negative, only the risk-sharing and credit-response terms on its right remain, which is the first condition directly. For that of (3.15), add the spillover channel $\varphi\, E^{\text{all}}[\Pi_{St}\Delta_t^\tau]$ to the risk-sharing term; rearranging yields the second. □

Conditional averaging. Writing S's spillover channel with an income-conditional average default response, as (3.10) does, is exact only if Π_{St} is conditionally uncorrelated with that response given y_t, and two things break the conditional independence that would deliver this. R's default rule depends on b_{Rt} while S's discount factor depends on b_{St}, and the two debt stocks inherit a common income history, so they are dependent given contemporaneous income; and Π_{St} itself moves with R's default status, which is the object being differentiated. The implementation carries the same approximation, weighting $\Delta(y)$ by S's mass at (b_S, y). This is a different approximation from the one that conditioning on income removes, and it affects the presentation rather than the argument: Proposition A.1 signs the state-level integrand directly and never invokes it.

A.3 The decompositions below the fiscal threshold

Section 3 derives the four welfare decompositions under $\tau \geq \bar{\tau}(\theta)$, where Proposition 2.1 makes $\omega_t \geq 0$ in every state. This subsection carries the general case. Nothing here bears on the quantitative results, since every Pareto-improving policy in the calibrated economy lies above the threshold, but the below-threshold terms are what make the threshold worth clearing in the first place.

Below $\bar{\tau}(\theta)$ the EC runs a deficit in low aggregate-income states, and a defaulting R is shielded from its share of it while S absorbs the whole shortfall. The default event $\mathcal{D}$ must therefore be split

by the sign of the surplus,

$$\mathcal{D}^{+} = \{d_{Rt} = 1,\ \omega_t \geq 0\}, \qquad \mathcal{D}^{-} = \{d_{Rt} = 1,\ \omega_t < 0\},$$

with $E^{\mathcal{D}} = E^{\mathcal{D}^{+}} + E^{\mathcal{D}^{-}}$ in the notation of (3.6). The fiscal threshold is exactly the value of τ above which $\mathcal{D}^{-}$ is empty, which is why the main text can work with the two-event partition.

Common debt. In $\mathcal{D}^{-}$ states R's EC receipt $\max\{\frac{1}{2}\omega_t, 0\}$ evaluates to zero and is insensitive to θ, so $\partial c_{Rt}/\partial\theta = 0$ rather than $-4(1-\bar{q})$: those periods drop out of R's interest cost, and (3.8) becomes

$$\frac{1}{u_R'(c_{R0})}\frac{d\mathcal{W}_R}{d\theta} \;=\; 4 \;-\; 4(1-\bar{q})\Big(E^{\mathcal{N}}[\Pi_{Rt}] + E^{\mathcal{D}^{+}}[\Pi_{Rt}]\Big) \;+\; E^{\mathcal{N}}\left[\Pi_{Rt}\, b_{Rt}'\,\frac{\partial q_R}{\partial\theta}\right]. \tag{A.8}$$

The waiver is a gain to R, so the impatience arbitrage of Section 3.2 is if anything larger below the threshold. For S the same states carry both $\frac{1}{2}\omega_t$ and $\min\{\frac{1}{2}\omega_t, 0\} = \frac{1}{2}\omega_t$, so the total EC term is ω_t and $\partial c_{St}/\partial\theta = -8(1-\bar{q})$ there. Equation (3.10) gains a *fiscal-exposure* term and its spillover channel picks up the avoided deficit share:

$$\begin{aligned}\frac{1}{u_S'(c_{S0})}\frac{d\mathcal{W}_S}{d\theta} \;=\; 4 \;-\; 4(1-\bar{q})E^{\text{all}}[\Pi_{St}] \;\underbrace{-\,4(1-\bar{q})E^{\mathcal{D}^{-}}[\Pi_{St}]}_{\text{fiscal exposure}} \\ +\; E^{\text{all}}\big[\Pi_{St}\,\Delta_t\big(\varphi + \max\{-\tfrac{\omega_t}{2}, 0\}\big)\big]\,.\end{aligned} \tag{A.9}$$

The fiscal-exposure term is non-positive and grows with θ at rate $4(1-\bar{q})$ per unit, so an untaxed program loads its entire deficit onto the safe country.

Common taxation. In $\mathcal{D}^{-}$ states R's EC receipt is fixed at zero regardless of τ, so $\partial c_{Rt}/\partial\tau = -y_{Rt}$: a higher tax reduces R's after-tax income with no offsetting rebate, which makes default marginally more costly. Using $-y_{Rt} = \frac{1}{2}(y_{St} - y_{Rt}) - \frac{1}{2}(y_{Rt} + y_{St})$, equation (3.13) gains a *default–EC tension* term:

$$\frac{1}{u_R'(c_{R0})}\frac{d\mathcal{W}_R}{d\tau} \;=\; E^{\text{all}}\left[\Pi_{Rt}\,\frac{y_{St} - y_{Rt}}{2}\right] \;\underbrace{-\,E^{\mathcal{D}^{-}}\left[\Pi_{Rt}\,\frac{y_{Rt} + y_{St}}{2}\right]}_{\text{default–EC tension}} \;+\; E^{\mathcal{N}}\left[\Pi_{Rt}\, b_{Rt}'\,\frac{\partial q_R}{\partial\tau}\right]. \tag{A.10}$$

It is non-positive: in these deficit-default states a defaulting R pays a higher τ without receiving a marginal rebate, so the perverse interaction of Section 2.1 imposes a pure cost precisely in distress. Symmetrically, in those states S bears the EC term ω_t in full, giving $\partial c_{St}/\partial\tau = y_{Rt} = \frac{1}{2}(y_{Rt} -$

$y_{St}) + \frac{1}{2}(y_{Rt} + y_{St})$, so (3.15) gains a *fiscal-exposure-reduction* term:

$$\frac{1}{u'_S(c_{S0})}\,\frac{d\mathcal{W}_S}{d\tau} \;=\; E^{\text{all}}\!\left[\Pi_{St}\,\frac{y_{Rt}-y_{St}}{2}\right] + \underbrace{E^{\text{all}}\!\left[\Pi_{St}\,\Delta_t^{\tau}\big(\varphi+\max\{-\tfrac{\omega_t}{2},0\}\big)\right]}_{\text{spillover channel}} + \underbrace{E^{\mathcal{D}^-}\!\left[\Pi_{St}\,\frac{y_{Rt}+y_{St}}{2}\right]}_{\text{fiscal-exposure reduction}}. \tag{A.11}$$

The fiscal-exposure-reduction term is non-negative, since a higher τ raises ω_t and shrinks the deficit S absorbs alone. The spillover channel picks up the avoided deficit share exactly as in (A.9): what S escapes when R repays is φ plus, in $\mathcal{D}^-$ states, the shortfall it would otherwise have absorbed alone, and that is the same object whichever instrument moves the default frequency.

Both threshold-contingent terms collapse to zero exactly at $\bar{\tau}(\theta)$, so the decomposition reproduces the fiscal threshold from the welfare side, and (A.8)–(A.11) reduce to their main-text counterparts. Below the threshold the flow budgets create a force pushing the θ credit response in an adverse direction, and the argument runs through the value functions rather than through the flow budget alone. Raising θ lowers R's repayment-branch resources by $4(1-\bar{q})$ in every period, but lowers its default-branch resources by $4(1-\bar{q})$ only in surplus states and by nothing in $\mathcal{D}^-$ states, where the contribution is waived. The repayment payoff therefore falls weakly more than the default payoff at every (b, y). That orders the Bellman-and-pricing operator in θ: evaluated at any fixed pair of continuation values and any fixed price schedule, a larger θ maps a given $V_R^c - V_R^d$ to a weakly smaller one, hence, the operator being monotone (Aguiar and Amador, 2019), to a weakly larger default set and a weakly lower price.

Propagating that ordering to the fixed point is a further step, and I do not take it. It would require the repayment branch to keep falling by weakly more at every iterate, which under partial reassumption of debt at re-entry does not follow immediately: a lower continuation value feeds back into the default branch through the re-entry term $\beta\alpha\,E[V_R((1-h)b, y')]$, so both branches fall and the ordering has to be shown to survive their difference. I therefore record $\partial q_R/\partial\theta \leq 0$ below the threshold as a property of the calibrated equilibrium rather than as a theorem, and it is measured directly there: holding $\tau = 0$, so that the whole θ range sits below $\bar{\tau}(\theta)$, raising θ from 0 to 0.5 raises R's annual default frequency from 1.09 to 1.15 percent and widens its spread from 1.02 to 1.06 percent (Section 4.2). The sign is in any case weak rather than strict, because the default set is a step correspondence in θ: over any range of program sizes that leaves it unchanged the derivative is exactly zero, which is also why the measured channel is flat above $\theta \approx 0.18$. That is the sense in which an untaxed common-debt program is worse for R's credit than a pure levy would be.

A.4 Alternative welfare measures

The headline welfare measure of Section 3.1, equation (3.2), is transition-inclusive: it integrates the new-regime value function $V_i(\cdot;\theta,\tau)$ against the *pre-reform* distribution $\mu^*(\cdot;0,0)$, and therefore

measures the change in lifetime utility of the agents alive when the program is enacted. This subsection reports two alternatives. They answer different questions and the paper keeps them apart:

	distribution	date-0 grant	question answered
$\mathcal{W}_i$	pre-reform $\mu^*(\cdot; 0, 0)$	yes	is adopting the reform good for those alive now?
$\mathcal{W}_i^{\text{own}}$	own $\mu^*(\cdot; \theta, \tau)$	yes	is that answer an artifact of the initial distribution?
$\mathcal{W}_i^{\text{ss}}$	own $\mu^*(\cdot; \theta, \tau)$, full support	no	what is the mature arrangement worth, once the transfer is history?

Robustness to the initial distribution. The first alternative weights the same date-0 value by each policy's own invariant distribution,

$$\mathcal{W}_i^{\text{own}}(\theta, \tau) \;\equiv\; \int V_i(b_i - 4\theta,\; y;\, \theta, \tau)\, d\mu^*(b_i, y;\, \theta, \tau). \tag{A.12}$$

This is the same launch experiment as (3.2), started from a different initial condition: the debt distribution the policy itself induces, rather than the one the economy actually inherits. Everything else is held identical. The date-0 grant still enters as the state shift $b_i - 4\theta$, the date-0 default branch of footnote 7 is still applied to R, and the distribution is still conditioned on R having market access and renormalized, jointly for both countries. The two measures differ *only* in the distribution against which that same date-0 value is integrated, and they coincide at the baseline $(\theta, \tau) = (0, 0)$, where the own distribution is the pre-reform distribution. It is a robustness check on the choice of initial distribution and nothing more; in particular it is not a stationary comparison, since it hands every sampled state a launch transfer.

The transition-inclusive object is the appropriate one for evaluating a reform from the status quo. It is the discrete consumption-equivalent analog of the marginal lifetime-utility perturbation of Dávila and Schaab (2025), which differentiates lifetime utility from a predetermined initial state and treats steady-state gains as the long-run limit of a full transition path, and it follows Dávila and Weymuller (2017) in valuing the welfare change at an initial date 0 starting from the no-program ergodic state. It is also internally consistent with the path-based decompositions of Section 3, which sum flow effects along the equilibrium path from date 0 rather than across stationary distributions.

Because the two share an integrand, the wedge between them is a single object,

$$\mathcal{W}_i(\theta, \tau) - \mathcal{W}_i^{\text{own}}(\theta, \tau) \;=\; \int V_i(b_i - 4\theta,\; y;\, \theta, \tau)\, d\big[\mu^*(b_i, y;\, 0, 0) - \mu^*(b_i, y;\, \theta, \tau)\big], \tag{A.13}$$

the new-regime value integrated against the displacement of the invariant distribution that the

program induces, and nothing else. It is large only when the reform relocates that distribution substantially *and* V_i varies sharply across the states where mass moves. Neither holds here. There is no aggregate capital to accumulate and income is exogenous and unchanged by the program, so the only route by which the program can move μ^* at all is through the equilibrium policy functions: an invariant distribution has no memory of an initial condition, so the date-0 grant washes out of it by construction and enters (A.13) only through the integrand. What moves the policy functions is the per-period fee $4(1-\bar{q})\theta$ and the tax τ, and the fee is worth 0.006 of one quarter's output per unit of θ. For S, whose savings problem sits at the $\beta_S = q_S$ knife-edge, the invariant law of a driftless random walk between reflecting barriers is independent of where it starts, and the distribution is almost unmoved.

Figure A.11 reports $\mathcal{W}_i^{\text{own}}$. The two measures are quantitatively close across the whole (θ, τ) grid: $\Delta\mathcal{W}_R$ runs to 10.69 percentage points against 10.42 under the headline measure, $\Delta\mathcal{W}_S$ agrees to three decimals over its entire range, and the Pareto region loses two of its 78 cells. What that establishes is precisely what (A.13) isolates, that the displacement of the invariant distribution is small, and not that transition-inclusive and long-run welfare coincide in general. The displacement would be first-order in a model with a slowly accumulated aggregate such as physical capital, or with overlapping generations, where current cohorts bear transition costs that a stationary comparison discards entirely.

The mature regime. The second alternative is the conventional stationary comparison. It drops the grant and integrates over the whole invariant support,

$$\mathcal{W}_i^{\text{ss}}(\theta, \tau) \;\equiv\; \int V_i(b_i,\; y;\, \theta, \tau)\, d\mu^*(b_i, y;\, \theta, \tau), \tag{A.14}$$

with no launch conditioning, since there is no launch date to condition on: R's excluded pool and S's $d_R = 1$ regime are both included. Both departures from (A.12) are forced by the same fact. The program's date-0 transfer is a one-time stock, an invariant distribution has forgotten it, and both countries have dissipated it: R because it is impatient and its ergodic debt is limit-determined, S because the invariant law of its reflected random walk does not depend on its starting point. So (A.14) prices the perpetual fee and the ongoing tax while excluding the resources they fund. It is not an alternative assessment of adopting the program, and it should not be read as one. It answers what a member faces in the mature regime, once the transfer is history.

Read that way it is informative, and it is adverse. Figure A.12 reports it: $\Delta\mathcal{W}_R^{\text{ss}}$ runs over $[-0.355,\, +0.010]$ percentage points and $\Delta\mathcal{W}_S^{\text{ss}}$ over $[-0.332,\, +0.006]$, and both are strictly negative at every $\theta > 0$. The 18 strictly Pareto-positive cells all lie in the $\theta = 0$ row, where the tax provides insurance without common debt. The two lowest-tax positive cells, at $\tau = 0.002$ and 0.004, remain at the numerical-noise scale, but from $\tau = 0.005$ through 0.020 both countries gain by more than

0.001 percentage points, so the insurance-only strip is not a coarse-grid artifact. The magnitude away from that strip is no accident: at $(\theta, \tau) = (0.5, 0)$ the measured figure of -0.32 percent for R sits against a per-period fee of $4(1 - \bar{q})\theta = 0.0030$, which is 0.31 percent of R's ergodic mean consumption. To a first approximation the mature-regime measure is minus the present value of the EC fee, which is exactly what Section 3.2 predicts: the scale of the program enters the recursive equilibrium through one term only, the per-period funding cost, so once the grant is outside the measure there is nothing else for θ to contribute.

Two things follow, and they pull in opposite directions. The first is a caveat on the headline result: R's gain is front-loaded by construction, so the program is not one a member would renew, and the paper claims no more than that adopting it is a Pareto improvement from the status quo. The second is a quantitative case for the institutional assumption of Section 2.1. Non-defaultability of the common tax is doing real work rather than serving as a convenience: a member that has spent its grant and faces the perpetuity strictly prefers to stop paying, and Figure A.12 is the size of that temptation. The one component that survives into the mature regime is the tax, whose Pareto-improving cells sit at $\theta = 0$: the insurance leg of the arrangement is self-sustaining in a way the mutualization leg is not.

A.5 Further results on common public goods

This subsection develops the public-good extension of Section 5.1 in full and collects its proof. Two ingredients carry over: period utility includes the additive term $\vartheta \log G_t$ of (5.1), and the EC devotes $\xi \geq 0$ per period to the common good out of its own budget (5.2), so each country's per-period transfer falls by $\xi/2$. The third ingredient is introduced here, and it is the one that generates the inefficiency the EC corrects: the countries also provide the good themselves, and they do so without coordinating.

Each country i maintains a local public-good authority that chooses $g_{it} \geq 0$ *after* the bond-issuance and default decisions are made, so local provision is financed out of after-debt net income and does not interact with default incentives directly. Let x_{it} denote the resources implied by the relevant budget constraint (2.6)–(2.9), with $\omega_t(\xi)$ replacing ω_t, before local public-good spending. Local provision is financed one-for-one out of these resources:

$$c_{it} + g_{it} \;=\; x_{it}, \qquad g_{it} \geq 0. \tag{A.15}$$

Thus $c_{it} = x_{it} - g_{it}$ is private consumption net of local public-good expenditure. Total provision is

$$G_t \;=\; g_{Rt} + g_{St} + \xi. \tag{A.16}$$

Nash equilibrium and underprovision. Taking $g_{-i,t}$, ξ, and x_{it} as given, country i chooses local provision by solving

$$\max_{g_{it}\geq 0} \; u_i(x_{it} - g_{it}) + \vartheta \log(g_{it} + g_{-i,t} + \xi).$$

The first-order condition holds with complementary slackness,

$$\frac{\vartheta}{G_t^N} \;\leq\; u_i'(c_{it}^N), \qquad \text{with equality whenever } g_{it}^N > 0, \tag{A.17}$$

so a contributing country equates the *private* marginal benefit of the public good to its marginal utility of consumption. The inequality matters: a single total provision G_t^N cannot satisfy both equalities unless $u_R'(c_{Rt}^N) = u_S'(c_{St}^N)$, so with unequal consumption the country with the higher marginal utility is generically at the corner $g_{it}^N = 0$ and only the other contributes, as in the classical private-provision result. A utilitarian social planner who internalizes both countries' benefits instead sets, for a contributing country,

$$\frac{2\vartheta}{G_t^{SP}} \;=\; u_i'(c_{it}^{SP}), \tag{A.18}$$

yielding $G^{SP} > G^N$: the standard fiscal-commons underprovision result (see Besley and Coate (2003) for the canonical analysis of centralized versus decentralized public-good provision). For symmetric countries under log-log preferences both conditions are homogeneous and the ratio simplifies to

$$\frac{G^{SP}}{G^N} \;=\; \frac{2+\vartheta}{1+\vartheta} \;\approx\; 1.91 \quad (\vartheta = 0.1), \tag{A.19}$$

confirming near-doubling of the socially optimal level relative to Nash provision. Equation (A.19) is an off-model illustration: (2.3) excludes $\gamma_i = 1$, so log-log preferences are used here only to obtain a closed form, and the calibrated economy sets $\gamma_i = 2$.

Welfare as a function of ξ. I evaluate welfare holding the recursive equilibrium (default decisions, debt policy functions, bond prices) fixed at the $\xi = 0$ solution. Holding the equilibrium fixed is a partial-equilibrium device, and it carries a second and more consequential restriction: each country's *partner* in the local provision game is held fixed as well. The paragraph *Scope of the fixed-partner calculation* below states what that restriction costs and why the envelope theorem does not cover it.

Under this assumption, an increase in ξ has two effects on country i: (i) private consumption falls by $\xi/2$ per period, because the EC surplus distributed to each country is reduced by $\xi/2$; and (ii) public-good provision rises by ξ. Country i's welfare at $\xi > 0$ is therefore

$$\mathcal{W}_i(\xi) \;=\; \frac{1}{1-\beta_i} E_\pi\Big[u_i\Big(c_{it} - \tfrac{\xi}{2}\Big) + \vartheta \log(G_t^N + \xi)\Big], \tag{A.20}$$

where c_{it} and G_t^N are the equilibrium consumption and Nash public-good level at $\xi = 0$. Differentiating with respect to ξ:

$$\frac{d\mathcal{W}_i}{d\xi} \;=\; \frac{1}{1-\beta_i}\, E_\pi\left[-\frac{1}{2}\, u_i'\Big(c_{it} - \tfrac{\xi}{2}\Big) + \frac{\vartheta}{G_t^N + \xi}\right]. \tag{A.21}$$

The two terms in (A.21) are the marginal public-good benefit and half the marginal utility of the private consumption it costs, and the race between them fixes the shape of $\mathcal{W}_i$. The first falls as total provision grows; the second rises as private consumption is squeezed. At $\xi = 0$ the Nash first-order condition (A.17) makes the two terms ϑ/G_t^N and $\frac{1}{2}u_i'(c_{it})$ differ by exactly $\frac{1}{2}u_i'(c_{it}) > 0$, the *inframarginal* gain: country i receives a public good worth ϑ/G^N per unit but pays only half its cost. At the other end, as ξ approaches twice the lowest ergodic consumption the marginal utility of what is foregone diverges. The next proposition turns those two limits into a unique interior optimum.

Figure A.1: Welfare effect of EC public-good expenditure ξ

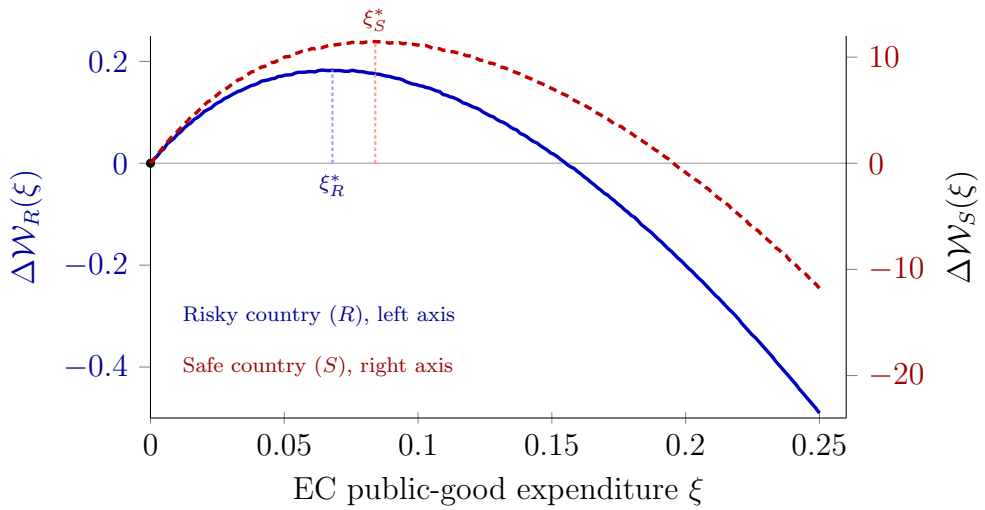

The figure shows how each country's lifetime welfare responds to the share ξ of the common budget spent on a public good, holding the recursive equilibrium fixed at the $\xi = 0$ allocation. It is a deterministic-steady-state evaluation at $(\theta, \tau) = (0.1, \bar{\tau}(0.1))$ with $\vartheta = 0.1$, $\gamma = 2$, $\bar{c}_R = 0.95$, $\bar{c}_S = 1.0$ and $G^N = 0.1$. Both curves are *lifetime* welfare changes, the per-period gain capitalized at each country's own discount factor as in (A.20). Because a permanent flow is worth $1/(1-\beta_S) \approx 668$ to S against $1/(1-\beta_R) \approx 14$ to R, the two are plotted on separate vertical scales, blue on the left for R and red on the right for S, sharing a common zero; the shapes and the optima are comparable across curves, the heights are not. Each country's welfare is hump-shaped in ξ (Proposition A.3): the initial gain from Nash underprovision is eventually dominated by the consumption cost. The dotted verticals mark the country-specific optima ξ_R^* and ξ_S^*. The curves are drawn from the closed-form expression (A.20) at the ergodic mean consumptions of the calibrated model rather than from a re-solved equilibrium, so they locate the optima under the fixed-equilibrium approximation of Proposition A.3 and not under the full solution of Section 5.1.

Proposition A.3 (Optimal EC Public-Good Provision)**.** *Holding the recursive equilibrium (default decisions, debt policy functions, bond prices) fixed at the $\xi = 0$ allocation, suppose that in every ergodic state country i values the first unit of centrally provided public good above the half unit of*

private consumption it costs,

$$\frac{\vartheta}{G_t^N} \;\geq\; \tfrac{1}{2}\, u_i'(c_{it}^N), \tag{A.22}$$

with strict inequality on a set of positive measure. Condition (A.22) *holds with slack* $\frac{1}{2}u_i'$ *for a contributing country, by* (A.17)*, and is the natural requirement for a country at the corner: the EC's implicit matching rate of one half is what can make provision worthwhile for a country that would not provide on its own. Then for each country* $i \in \{R, S\}$*, the welfare function* $\mathcal{W}_i(\xi)$ *is strictly concave with a unique interior maximum* $\xi_i^* > 0$ *characterized by*

$$E_\pi\left[\frac{\vartheta}{G_t^N + \xi_i^*}\right] \;=\; \frac{1}{2}\, E_\pi\left[u_i'\left(c_{it} - \frac{\xi_i^*}{2}\right)\right]. \tag{A.23}$$

In the deterministic steady state, this reduces to

$$\frac{\vartheta}{G^N + \xi_i^*} \;=\; \frac{1}{2}\left(\bar{c}_i - \frac{\xi_i^*}{2}\right)^{-\gamma}, \tag{A.24}$$

which under $\gamma = 1$ *yields the closed form* $\xi_i^* = (2\vartheta\bar{c}_i - G^N)/(1+\vartheta)$*, again an off-model illustration since* (2.3) *excludes* $\gamma_i = 1$*.*

Proof. From (A.21), the marginal welfare of country i at ξ is

$$\frac{d\mathcal{W}_i}{d\xi} = \frac{1}{1-\beta_i}\, E_\pi\left[\frac{\vartheta}{G_t^N + \xi} - \frac{1}{2}\, u_i'\left(c_{it} - \frac{\xi}{2}\right)\right].$$

Step 1 (positive at $\xi = 0$*).* At $\xi = 0$ the integrand is $\vartheta/G_t^N - \frac{1}{2}u_i'(c_{it})$, which is non-negative in every state by (A.22) and strictly positive on a set of positive measure; for a contributing country (A.17) holds with equality and the integrand equals $\frac{1}{2}u_i'(c_{it}) > 0$ pointwise. Hence $d\mathcal{W}_i/d\xi|_{\xi=0} > 0$.

Step 2 (strictly decreasing). Write $m_i(\xi; s) = \vartheta/(G^N(s) + \xi) - \frac{1}{2}u_i'(c_i(s) - \xi/2)$ for each ergodic state s. The first term is strictly decreasing in ξ and the second is strictly increasing in magnitude (since $u_i'' < 0$ implies u_i' is increasing as its argument decreases). Hence $m_i(\xi; s)$ is strictly decreasing in ξ for every s, and so is $d\mathcal{W}_i/d\xi = (1-\beta_i)^{-1}E_\pi[m_i]$.

Step 3 (tends to $-\infty$*).* As $\xi \to 2\min_s c_i(s)$, private consumption $c_i - \xi/2 \to 0$ and $u_i'(c_i - \xi/2) \to \infty$, so $d\mathcal{W}_i/d\xi \to -\infty$.

Step 4 (unique interior maximum). Since $d\mathcal{W}_i/d\xi$ is continuous, starts positive (Step 1), is strictly decreasing (Step 2), and tends to $-\infty$ (Step 3), by the intermediate value theorem there exists a unique $\xi_i^* > 0$ satisfying $d\mathcal{W}_i/d\xi = 0$, which is the global maximum. The first-order condition at ξ_i^* gives (A.23). The deterministic steady-state reduction (A.24) follows by replacing expectations with steady-state values. □

Remark A.1 (Utilitarian Optimum and the Samuelson Rule)**.** *The utilitarian optimum* ξ^* *maximizes* $\mathcal{W}_R(\xi) + \mathcal{W}_S(\xi)$*. Because* (A.20) *capitalizes each country's per-period gain at its own discount*

factor, the two countries enter the sum with weights $\omega_i \equiv 1/(1-\beta_i)$*, and the Samuelson condition of Samuelson (1954) reads*

$$(\omega_R + \omega_S)\, E_\pi\left[\frac{\vartheta}{G_t^N + \xi^*}\right] \;=\; \frac{1}{2}\, E_\pi\left[\omega_R\, u_R'\!\left(c_{Rt} - \frac{\xi^*}{2}\right) + \omega_S\, u_S'\!\left(c_{St} - \frac{\xi^*}{2}\right)\right], \tag{A.25}$$

equating the sum *of marginal public-good benefits to the* sum *of the marginal costs, each country bearing half. At the baseline calibration with* $(\theta, \tau) = (0.1,\, \bar{\tau}(0.1))$*,* $\bar{c}_R \approx 0.95$*,* $\bar{c}_S \approx 1.0$*, and* $G^N \approx 0.1$*, the deterministic steady-state approximation yields* $\xi_R^* \approx 0.068$*,* $\xi_S^* \approx 0.084$*, and* $\xi^* \approx 0.083$*. The utilitarian optimum sits almost on top of* ξ_S^* *because* $\omega_S/\omega_R \approx 47$ *under the knife-edge* $\beta_S = q_S$*: the patient safe country's preference dominates the sum. Dropping the weights, that is, adding the two per-period gains rather than the two lifetime values, would instead give* $\xi^* \approx 0.075$*.*

Scope of the fixed-partner calculation. Equation (A.20) holds each country's partner contribution fixed as ξ varies, and this is an approximation rather than an envelope result. The envelope theorem removes the first-order effect of country i's own re-optimization of g_{it}, because g_{it} was chosen optimally beforehand. It does not remove the effect of the *partner's* response, since $\partial u_i/\partial g_{-i,t} = \vartheta/G_t > 0$, and $g_{-i,t}$ falls when the EC provides.

The omission is not small, and its limiting case is instructive. Substituting $\hat{g}_{it} = g_{it} + \xi/2$ turns the local provision game into the $\xi = 0$ game subject to $\hat{g}_{it} \geq \xi/2$, because the EC funds ξ by withdrawing $\xi/2$ from each member. If both countries contribute strictly and that constraint is slack, the $\xi = 0$ profile survives unchanged: total provision, consumption and welfare are all invariant to ξ, exactly, and the program is neutral in the sense of Bergstrom et al. (1986). Centralized provision therefore does something only when at least one country sits at the corner $g_{it} = 0$ and cannot absorb the levy by cutting a contribution it is not making. That is the generic configuration here, since a single G_t^N cannot satisfy both first-order conditions at unequal marginal utilities, but it means the extension's gains rest on the corner rather than on non-rivalry alone.

With one contributor and one corner country, the contributor's marginal gain is $\frac{1}{2}u_j'(c_{jt})$ as in Step 1 of the proposition, but total provision rises by less than ξ: writing $A = \vartheta/G^2$ and $B = -u_j''$, the contributor's response gives $dG/d\xi = (B/2)/(A+B) < \frac{1}{2}$, which at the calibrated $\vartheta = 0.1$ and $G^N \approx 0.091$ is about 0.08. Equation (A.20) instead sets $dG/d\xi = 1$.

The direction in which that biases the surrogate optima is not the same for the two countries, and revealed preference settles it without any derivative. Holding both contributions frozen remains feasible for the contributor at every ξ, and the corner country stays at its corner as ξ rises, so the contributor can always attain the frozen allocation and its true welfare is weakly above it: (A.20) is a *lower* bound on the contributor's gain. The envelope derivative says the same thing more sharply. At any ξ at which the contributor still contributes, its own first-order condition (A.17) gives $\vartheta/G_t = u_j'(c_{jt})$, so the marginal welfare integrand of (A.21) is $\vartheta/G_t - \frac{1}{2}u_j'(c_{jt}) = \frac{1}{2}u_j'(c_{jt}) > 0$

in every such state, with the contributor's own response to ξ already netted out by the envelope theorem. A contributing country therefore has no interior optimum while it contributes, and its true optimum lies beyond the point at which it stops. For the corner country the comparison runs the other way. Its consumption path is the same under either calculation, while true provision $g_{jt} + \xi$ falls short of the frozen $G_t^N + \xi$ because the contributor retrenches, so (A.20) is an *upper* bound on its gain. The response $dG/d\xi < 1$ of the previous paragraph is thus the corner country's bound; for the contributor that same response is what turns its marginal gain into $\frac{1}{2}u_j'$. Figure A.1 accordingly overstates the case for centralized provision as the corner country sees it and understates it as the contributor does, and by (A.17) the contributor is the country with the *lower* marginal utility of consumption, which at the illustration's $\bar{c}_R < \bar{c}_S$ is the safe country. Which way the utilitarian ξ^* of Remark A.1 moves is therefore not signed by this argument, and neither is the lower of the two countries' preferred levels, which is the Pareto-relevant one: the correction raises the contributor's optimum and lowers the corner country's, and it is the corner country's that binds.

Two limits of the comparison are worth stating. It holds within a configuration of one contributor and one corner country, so it transfers directly to Figure A.1, which is a single deterministic state, but it does not aggregate into a bound on a country's welfare across the solved surface, where which country sits at the corner varies with the state. In states where both contribute the gap runs the other way again, since neutrality leaves true welfare invariant to ξ while the frozen calculation still rises at first. And nothing here relocates the optimum, only its direction of bias, so the statement in Section 5.1 that $\xi = 0.05$ lies below the benchmark optimum is a statement about the optima as Proposition A.3 computes them, not about their corrected values.

The numerical implementation makes a related but distinct approximation, described under *Computing local provision* below: it fixes the partner at the symmetric Nash contribution of a country facing mean resources $\bar{x} = 1$ at the prevailing ξ. Because that rule evaluates the partner at un-netted resources rather than at $\bar{x} - \xi/2$, it delivers a positive but small provision response, G rising from 0.0911 to 0.0955 as ξ goes from 0 to 0.05; netting the levy out of the partner's resources would restore exact neutrality.

What the reported surface turns on is not the sign of that response but the level at which the partner is pinned, and that has been probed directly. Re-solving twelve cells spanning the reported sign changes, $\theta \in \{0, 0.18, 0.54\}$ against $\tau \in \{0.0025, 0.0075, 0.0125, 0.015\}$, the largest of which reaches just past the top of the policy grid, at four values of the frozen partner contribution ranging over a factor of four, from 0.036 to 0.146 against the calibrated 0.073, and re-solving each column's own $(0, 0)$ benchmark so that each is a self-contained economy, moves $\Delta\mathcal{W}_S$ by at most 0.010 percentage points, six percent of the largest gain involved, and changes no cell's sign. The range is wide enough to straddle the regime distinction of the previous paragraph, since at the highest value a country facing mean resources is at the corner while at the other three it contributes. The results of Section 5.1 therefore do not turn on where that constant is pinned, and the three percent

question of whether to net the levy is immaterial a fortiori.

Figure A.1 illustrates the welfare change $\Delta\mathcal{W}_i(\xi) \equiv \mathcal{W}_i(\xi) - \mathcal{W}_i(0)$ as a function of ξ at $(\theta, \tau) = (0.1, \bar{\tau}(0.1))$. Both countries benefit from moderate public-good provision: at $\xi = 0$, the marginal public-good benefit exceeds the marginal consumption cost (Nash underprovision), but as ξ increases, the diminishing marginal utility of G and the rising marginal utility of foregone consumption produce a unique interior peak. The safe country S can sustain a slightly higher ξ because its higher average consumption makes the cost $\xi/2$ relatively less painful.

Computing local provision. An exact implementation of (A.17) requires, at every state in the VFI, a Nash fixed point over the pair $(g_{R,t}, g_{S,t})$: each country's optimal g_{it} depends on the other's through $G_t^N = g_{Rt} + g_{St} + \xi$, and both depend on the state-dependent resources x_{it}. The solver avoids this nested fixed point by holding the non-own component of G_t fixed at a state-independent level $\bar{G}^{-i} = \bar{g} + \xi$, where $\bar{g}$ is the symmetric Nash contribution of a country facing mean resources $\bar{x} = 1$ and the EC's own expenditure, that is the solution of $\vartheta/(2\bar{g}+\xi) = u_i'(\bar{x}-\bar{g})$. Country i's local problem then reduces to a one-dimensional choice $g_{it} \geq 0$ solving $\vartheta/(g_{it}+\bar{G}^{-i}) = u_i'(x_{it}-g_{it})$, whose left side falls and whose right side rises in g_{it}, so the root is unique. At the calibrated $\gamma_i = 2$ the condition is a quadratic and the root is available in closed form, which is what the solver uses; for a general CRRA curvature it is a one-dimensional monotone root-find rather than a closed form. It is evaluated pointwise inside the VFI, with the corner $g_{it} = 0$ imposed wherever it binds. Two errors are introduced: evaluating the partner's contribution at mean resources rather than averaging its policy over the ergodic distribution, and freezing it across states. Neither is an envelope-theorem casualty, for the reason given under *Scope of the fixed-partner calculation* above, and the second is the more interesting of the two, since a partner that responded to its own state would contribute less exactly when country i is poor, so that freezing $\bar{G}^{-i}$ removes a covariance that would otherwise raise ϑ/G_t in bad states. The direction in which that biases the (θ, τ) comparative statics is not signed here. What can be said is quantitative: varying the level at which the partner is frozen over a factor of four leaves every probed cell's welfare sign intact and moves $\Delta\mathcal{W}_S$ by at most six percent of the largest gain involved. A fully-nested Nash solver is straightforward to code but costs roughly an order of magnitude more runtime per grid point; the approximation keeps the extension tractable alongside the baseline.

A.6 Further results on convenience yield

This subsection develops the convenience-yield extension of Section 5.3 in full and collects the proofs of its three results. The single change to the baseline is that the EC bond rate depends on program scale, $r_J(\theta) = \max\{r_S + \eta - \nu\theta,\, \bar{r}_J\}$ as in (5.6), so the issue price $\bar{q}(\theta) = 1/(1 + r_J(\theta))$ is no longer a constant. Everything else is inherited: η is the initial illiquidity premium, ν the market-depth

elasticity, $\bar{r}_J \in (0, r_S)$ the long-run floor reached at $\theta^\dagger$, and η/ν the scale at which common debt prices exactly like the safe sovereign's.

All equations of Section 2 carry through with $\bar{q}(\theta)$ replacing $\bar{q}$. The EC surplus becomes

$$\omega_t(\theta) \;=\; \tau(y_{R,t} + y_{S,t}) - 8\big(1 - \bar{q}(\theta)\big)\theta \;=\; \tau(y_{R,t} + y_{S,t}) - \frac{8\, r_J(\theta)}{1 + r_J(\theta)}\,\theta. \tag{A.26}$$

For $\theta \geq \theta^\dagger$ the per-unit debt-service cost stabilizes at $\bar{r}_J/(1 + \bar{r}_J)$.

Under (5.6) the program is no longer cash-flow neutral for S. As in the baseline, θ indexes the cash the EC raises, so the date-0 grant stays at 4θ per country; what moves is the per-period service cost $8(1 - \bar{q}(\theta))\theta$ needed to raise it. Since S discounts that stream at its own rate q_S rather than at the EC's $\bar{q}(\theta)$, the net date-0 value of the program to S is $4\theta\big[1 - (1 - \bar{q}(\theta))/(1 - q_S)\big]$, which is negative while $r_J(\theta) > r_S$, zero at $\theta = \eta/\nu$, and positive once the convenience yield takes hold. The baseline of Proposition 3.1 is the knife-edge $\bar{q} \equiv q_S$ at which this expression vanishes identically.

Non-monotone fiscal threshold. Substituting (5.6) into Proposition 2.1's formula $\bar{\tau}(\theta) = 8(1 - \bar{q}(\theta))\theta/y^{min}$ gives

$$\bar{\tau}(\theta) \;=\; \begin{cases} \dfrac{8(r_S + \eta - \nu\theta)}{(1 + r_S + \eta - \nu\theta)\, y^{min}}\,\theta & \text{if } \theta < \theta^\dagger, \\[2ex] \dfrac{8\bar{r}_J}{(1 + \bar{r}_J)\, y^{min}}\,\theta & \text{if } \theta \geq \theta^\dagger. \end{cases} \tag{A.27}$$

Proposition A.4 (Non-Monotone Fiscal Threshold)**.** *Let $r_J(\theta) = \max\{r_S + \eta - \nu\theta,\, \bar{r}_J\}$ with $\eta, \nu > 0$ and $0 < \bar{r}_J < r_S$, and define*

$$\theta^{**} \;=\; \frac{(1 + r_S + \eta) \;-\; \sqrt{1 + r_S + \eta}}{\nu}, \qquad \theta^\dagger \;=\; \frac{r_S + \eta - \bar{r}_J}{\nu}. \tag{A.28}$$

The fiscal threshold $\bar{\tau}(\theta)$ in (A.27) *satisfies:*

(i) $\bar{\tau}'(\theta) > 0$ for $\theta < \theta^{**}$: the illiquidity premium dominates.

(ii) If $\theta^{**} < \theta^\dagger$, so that the interval is nonempty, $\bar{\tau}'(\theta) < 0$ for $\theta \in (\theta^{**}, \theta^\dagger)$: the convenience yield reduces per-unit costs faster than the bond stock grows.

(iii) For $\theta \geq \theta^\dagger$, the convenience-yield floor binds and $\bar{\tau}(\theta)$ grows linearly at slope $8\bar{r}_J/((1 + \bar{r}_J)y^{min})$. If $\theta^{**} < \theta^\dagger$, then $\theta^\dagger$ is a local minimum.

(iv) $\bar{\tau}(\theta)$ crosses the baseline threshold $8r_S\theta/((1 + r_S)y^{min})$ exactly at $\theta = \eta/\nu$: it exceeds the baseline on $(0, \eta/\nu)$, coincides with it at $\theta = \eta/\nu$ (where $r_J(\theta) = r_S$), and lies strictly below it for $\theta > \eta/\nu$.

Proof. Write $a \equiv r_S+\eta$. On $[0,\theta^\dagger)$ the floor is slack, so $\bar{q}(\theta) = 1/(1+a-\nu\theta)$ and $\bar{\tau}(\theta) = (8/y^{min})\,(a-\nu\theta)\theta/(1+a-\nu\theta)$. Setting $\bar{\tau}'(\theta) = 0$ and clearing denominators yields $\nu^2\theta^2 - 2\nu(1+a)\theta + a(1+a) = 0$, with roots $[(1+a) \pm \sqrt{1+a}]/\nu$. The larger root exceeds a/ν and lies outside the domain, so the only interior stationary point is the θ^{**} of (A.28), and the sign of the quadratic gives $\bar{\tau}' > 0$ below it and $\bar{\tau}' < 0$ on $(\theta^{**},\theta^\dagger)$ whenever that interval is nonempty, which is (i) and (ii). For (iii), $r_J = \bar{r}_J$ for $\theta \geq \theta^\dagger$, so $\bar{\tau}$ is linear there with slope $8\bar{r}_J/((1+\bar{r}_J)y^{min}) < 8r_S/((1+r_S)y^{min})$ since $\bar{r}_J < r_S$; if $\theta^{**} < \theta^\dagger$ the left derivative at $\theta^\dagger$ is negative and the right derivative positive, so $\theta^\dagger$ is a local minimum. For (iv), $\bar{\tau}(\theta)$ is $8\theta/y^{min}$ times $x/(1+x)$ evaluated at $x = r_J(\theta)$, and $x/(1+x)$ is strictly increasing in x, so $\bar{\tau}$ compares with the baseline threshold exactly as $r_J(\theta)$ compares with r_S: above on $(0,\eta/\nu)$, equal at η/ν, and below thereafter, including on $[\theta^\dagger,\theta^{\max}]$ where $r_J = \bar{r}_J < r_S$. Since $\theta^\dagger > \eta/\nu$, that crossing is the unique one on $\theta > 0$. □

Figure A.2: Variable convenience yield and the non-monotone fiscal threshold

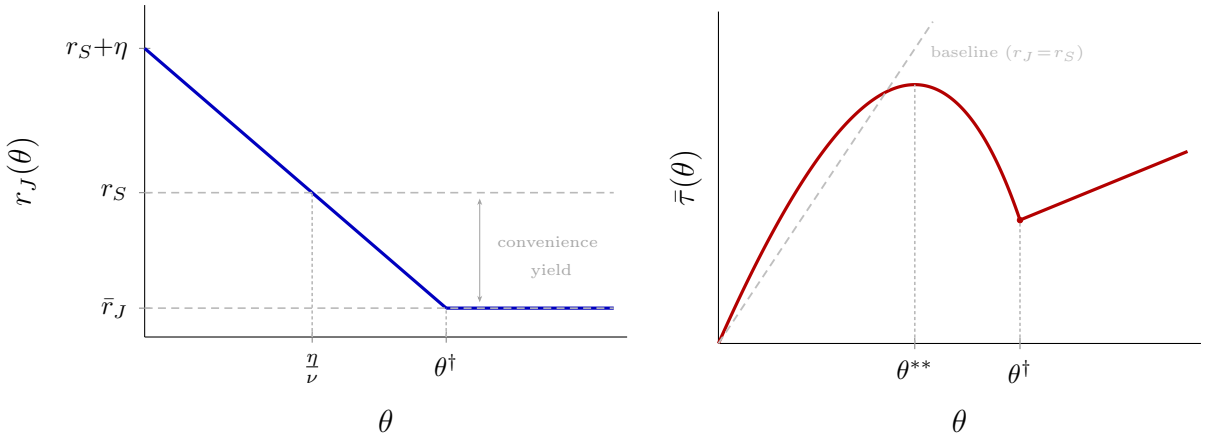

The figure shows the two objects that change when the price at which the common institution issues is allowed to depend on the scale of the program through (5.6). *Left*: the EC bond rate $r_J(\theta)$ is decreasing in the per-country transfer θ; it falls below r_S at $\theta = \eta/\nu$ and reaches the floor $\bar{r}_J$ at $\theta^\dagger$, where the convenience yield saturates. *Right*: the fiscal threshold $\bar{\tau}(\theta)$ is single-peaked on the nonbinding segment when $\theta^{**} < \theta^\dagger$, peaking at the tipping point θ^{**}. Once the floor binds, it grows linearly at the reduced slope $8\bar{r}_J/((1+\bar{r}_J)\,y^{min})$; in the case $\theta^{**} < \theta^\dagger$, $\theta^\dagger$ is a local minimum. The dashed line on the right panel shows the baseline threshold under the constant $r_J = r_S$ of Section 2, which the variable-yield threshold crosses once, at η/ν. The parameter values behind the two panels ($r_S+\eta = 1$, $\nu = 2$, $\bar{r}_J = 0.1$, $r_S = 0.5$, and y^{min} normalized to one) are purely illustrative and chosen to separate the thresholds visibly; they are far from the calibration of Table 4.1, the model is not re-solved under (5.6), and the figure is a qualitative guide to the shapes rather than a quantitative one.

Non-monotone safe-country welfare. The non-monotone fiscal threshold has a direct welfare counterpart. To see this, define the *fee factor* as a function of θ, of which the per-unit cost is $1-\bar{q}(\theta)$ and the per-country fee $4\Lambda(\theta)$:

$$\Lambda(\theta) \;\equiv\; \big(1-\bar{q}(\theta)\big)\,\theta \;=\; \frac{r_J(\theta)}{1+r_J(\theta)}\,\theta. \tag{A.29}$$

On $[0, \theta^{\dagger})$, Λ inherits the hump shape of $\bar{\tau}$ when the interior turning point is reached before the floor binds: it peaks near θ^{**} and, if $\theta^{**} < \theta^{\dagger}$, reaches a local minimum at $\theta^{\dagger}$. Beyond $\theta^{\dagger}$ it grows linearly at slope $\bar{r}_J/(1+\bar{r}_J)$.

Above the fiscal threshold, three forces shape $d\mathcal{W}_S/d\theta$: the *marginal cash grant*, fixed at 4 per unit of θ by the convention that θ indexes cash raised; the *marginal EC interest cost* $4\Lambda'(\theta)$, with $\Lambda'(\theta) = (1-\bar{q}(\theta)) - \theta\,\bar{q}'(\theta)$, which is strictly decreasing on $[0, \theta^{\dagger})$, jumps upward when the floor binds, and is constant at $\bar{r}_J/(1+\bar{r}_J)$ above it; and the *spillover benefit* $B(\theta) \geq 0$, non-decreasing in θ, which is maintained as an assumption here rather than derived. It is worth being explicit that the assumption is not inherited from the quantitative model: at the constant $r_J = r_S$ of the baseline, S's θ channel is measured slightly *negative* in Table 4.2, because a larger untaxed program worsens R's credit. Assumption (a) is therefore a restriction the variable-rate economy would have to satisfy, and which a scale-dependent funding cost could plausibly deliver, rather than a property carried over from Section 4.3. When η is large the marginal EC cost overwhelms the spillover benefit at small θ, while near $\theta^{\dagger}$ it has fallen to $\bar{r}_J/(1+\bar{r}_J)$ and the spillover benefit dominates, so $d\mathcal{W}_S/d\theta$ changes sign.

Proposition A.5 (Non-Monotone Safe-Country Welfare)**.** *Under the variable convenience yield* (5.6) *with floor* $\bar{r}_J \in (0, r_S)$*, fix a tax rate* τ *constant in* θ *and high enough to clear the threshold over the whole range,* $\tau \geq \max_{\theta \in [0,\theta^{\max}]} \bar{\tau}(\theta)$*, which exists because* $\bar{\tau}$ *is continuous on a compact interval, so that* (3.10) *applies throughout and* $d/d\theta$ *is a total derivative. Assume:*

(a) $B(\theta)$ *is non-decreasing, continuously differentiable and convex, with* $B(0) = 0$*. Convexity is used only for the uniqueness half of part (iii); existence of a trough does not need it.*

(a′) *The equilibrium-response terms that* (3.10) *carries beyond the grant, the fee and the spillover benefit, namely the movement of the SDF weights* Π_{St} *and of the invariant distribution as* θ *varies, are collected in a continuous residual* $\mathcal{R}(\theta)$ *on* $[0, \theta^{\max}]$ *satisfying*

$$\|\mathcal{R}\|_{\infty} \;\equiv\; \sup_{\theta \in [0,\theta^{\max}]} \big|\mathcal{R}(\theta)\big| \;<\; \min\big\{\textit{slack in } (\text{A.31}),\ \textit{slack in } (\text{A.32})\big\}, \tag{A.30}$$

and non-decreasing on $[0, \theta^{\dagger})$*.*[30]

[30] Both halves are assumptions, not consequences, and they are why (b) and (c) are stated as strict inequalities. The bound (A.30) is two-sided because the two endpoint conditions are exposed to residuals of opposite sign, a large positive one at 0^{+} and a large negative one at $\theta^{\dagger}$; it delivers parts (i) and (ii) and hence the *existence* of a trough. Monotonicity is used only for the uniqueness half of part (iii), where it plays the role convexity of B plays in (a), since a uniformly small residual may still oscillate near the zero of the non-residual expression and produce several crossings. It can be replaced by any condition making the complete marginal expression strictly increasing below $\theta^{\dagger}$, for instance differentiability with $\|\mathcal{R}'\|_{\infty}$ below the infimum over $[0, \theta^{\dagger})$ of the derivative of the non-residual part of (A.33). Here and in (b) and (c), $E^{\text{all}}[\Pi_{St}]$ is the SDF weight at a fixed reference policy, a constant rather than a function of θ, and its movement with the program is one of the terms $\mathcal{R}$ collects.

(b) At $\theta = 0^+$, the marginal EC cost exceeds the cash grant plus spillover benefit:

$$4\Lambda'(0^+)\, E^{all}[\Pi_{St}] \;>\; 4 + B'(0^+). \tag{A.31}$$

(c) At $\theta = \theta^\dagger$, the spillover benefit overcomes the residual EC cost:

$$B'(\theta^\dagger) + 4 \;>\; \frac{4\,\bar{r}_J}{1+\bar{r}_J}\, E^{all}[\Pi_{St}]. \tag{A.32}$$

Then $\mathcal{W}_S(\theta, \tau)$ is non-monotone in θ:

(i) $d\mathcal{W}_S/d\theta < 0$ near $\theta = 0$.

(ii) $d\mathcal{W}_S/d\theta > 0$ on a neighborhood of $\theta^\dagger$ and everywhere above it.

(iii) There is a unique *$\theta^v \in (0, \theta^\dagger)$ with $d\mathcal{W}_S/d\theta\big|_{\theta^v} = 0$, and $d\mathcal{W}_S/d\theta < 0$ on $(0, \theta^v)$ and > 0 on $(\theta^v, \theta^\dagger)$: a single welfare trough.*

Proof. Above the fiscal threshold ($\tau \geq \bar{\tau}$), S's welfare decomposition (3.10) gives, under the variable convenience yield,

$$\frac{1}{u'_S(c_{S0})}\,\frac{d\mathcal{W}_S}{d\theta} \;=\; \underbrace{4 - 4\Lambda'(\theta)\, E^{\text{all}}[\Pi_{St}] + B'(\theta)}_{\equiv\; F(\theta)} \;+\; \mathcal{R}(\theta), \tag{A.33}$$

where 4 is the marginal cash grant and $4\Lambda'(\theta) = 4\big[(1 - \bar{q}(\theta)) - \theta\,\bar{q}'(\theta)\big]$ is the marginal EC fee, both per unit of θ, and $\mathcal{R}$ collects the equilibrium-response terms as in assumption (a′). Since $u'_S(c_{S0}) > 0$, the sign of $d\mathcal{W}_S/d\theta$ is the sign of $F + \mathcal{R}$ throughout, and it is that sum, not F alone, that the three parts must sign. The convenience yield enters through Λ' alone: a deeper market lowers the service cost of the cash the EC raises, and $-\theta\bar{q}'(\theta)$ is the saving on the inframarginal units.

The marginal fee in closed form. Write $a \equiv r_S + \eta$ and $u(\theta) \equiv 1 + r_J(\theta)$. On $[0, \theta^\dagger)$ the floor is slack, $u = 1 + a - \nu\theta$ and $\bar{q} = 1/u$, so $\bar{q}' = \nu/u^2$ and, using $\nu\theta = (1 + a) - u$,

$$\Lambda'(\theta) \;=\; \frac{u-1}{u} - \frac{\nu\theta}{u^2} \;=\; 1 - \frac{1+a}{u(\theta)^2}. \tag{A.34}$$

Since u is strictly decreasing in θ and strictly positive, $(1 + a)/u^2$ is strictly increasing, so Λ' is *strictly decreasing* on $[0, \theta^\dagger)$: the scale economy on inframarginal units grows faster than the per-unit cost falls. At $\theta^\dagger$ the floor binds, $\bar{q}'$ drops to zero, and Λ' jumps *up* from $1 - (1 + a)/(1 + \bar{r}_J)^2$ to $\bar{r}_J/(1 + \bar{r}_J) = 1 - 1/(1 + \bar{r}_J)$, the jump being upward because $a > \bar{r}_J$. Above $\theta^\dagger$ it is constant.

Part (i): at $\theta = 0^+$, (A.34) gives $\Lambda'(0^+) = 1 - 1/(1 + a) = (r_S + \eta)/(1 + r_S + \eta) > 0$, which strictly exceeds $1 - q_S$ for $\eta > 0$. Condition (A.31) then ensures $4\Lambda'(0^+)\, E^{\text{all}}[\Pi_{St}] > 4 + B'(0^+)$, that

is $F(0^+) < 0$, by a margin equal to the slack in (A.31). The bound (A.30) keeps $|\mathcal{R}(0^+)|$ strictly below that margin, so $F(0^+) + \mathcal{R}(0^+) < 0$ and $d\mathcal{W}_S/d\theta < 0$. By continuity of F and of $\mathcal{R}$, the inequality holds in a neighborhood of $\theta = 0$.

Part (ii): for $\theta \geq \theta^\dagger$ (floor binds), $\bar{q}(\theta) = 1/(1+\bar{r}_J)$ is constant so $\bar{q}'(\theta) = 0$ and the marginal fee is $\Lambda' = 1 - \bar{q} = \bar{r}_J/(1+\bar{r}_J)$. There $F(\theta) = 4 - 4(\bar{r}_J/(1+\bar{r}_J))\, E^{\text{all}}[\Pi_{St}] + B'(\theta)$, which is strictly positive at $\theta^\dagger$ by condition (A.32) and, B' being non-decreasing, remains so for every $\theta > \theta^\dagger$; the margin is the slack in (A.32), which (A.30) again strictly exceeds in absolute value, so $F + \mathcal{R} > 0$ there. This is the half of the two-sided bound that protects against a large *negative* residual, the opposite sign from the one that matters in part (i). Immediately to the *left* of $\theta^\dagger$ the marginal fee is strictly smaller, by the upward jump established above, so the left-limit of F at $\theta^\dagger$ strictly exceeds its value at $\theta^\dagger$ and $F + \mathcal{R}$ is a fortiori positive. Hence $d\mathcal{W}_S/d\theta > 0$ on a two-sided neighborhood of $\theta^\dagger$ and everywhere above it.

Part (iii): on $[0, \theta^\dagger)$, B' is non-decreasing by convexity in (a) and $-\Lambda'$ is strictly increasing by (A.34), so F is strictly increasing there. Strict monotonicity of F alone does not deliver a single crossing, because what changes sign is $F + \mathcal{R}$: a residual bounded by (A.30) may still oscillate, and an arbitrarily small oscillation in a neighborhood of the zero of F produces several zeros of the sum. This is where the second half of (a′) is used. With $\mathcal{R}$ non-decreasing on $[0, \theta^\dagger)$, the sum $F + \mathcal{R}$ is strictly increasing there; it is negative at 0^+ by Part (i) and positive as $\theta \uparrow \theta^\dagger$ by Part (ii), so by the intermediate value theorem it has exactly one zero θ^v in $(0, \theta^\dagger)$, is negative below it and positive above it. Under the alternative form of (a′) the same conclusion follows from $(F + \mathcal{R})' \geq F' - \|\mathcal{R}'\|_\infty > 0$. □

Both conditions have a transparent reading once $E^{\text{all}}[\Pi_{St}] = 1/(1 - q_S)$ is substituted, which the knife-edge $\beta_S = q_S$ delivers. The substitution needs S's Euler equation to hold with equality, so it is exact only where S's debt choice is interior, and the reflecting band of Section 2 binds in a non-negligible share of periods in the calibrated economy. That is one more reason this subsection is developed as theory rather than re-solved. Condition (A.31) then reduces to $(1 - \bar{q}(0))/(1 - q_S) > 1 + B'(0^+)/4$, which holds for any $\eta > 0$ large enough relative to the spillover benefit, since the illiquidity premium makes the EC's service cost exceed what S would pay on its own account. Condition (A.32) reduces to $B'(\theta^\dagger)/4 + 1 > \big[\bar{r}_J/(1+\bar{r}_J)\big]\big/\big[r_S/(1+r_S)\big]$, whose right-hand side is below one whenever $\bar{r}_J < r_S$: at the floor the EC's cost of funds lies strictly below S's own, so the program pays for itself out of the grant alone. Since $B'(\theta^\dagger) \geq 0$ by (a), the left-hand side is at least one, and the condition is implied by $\bar{r}_J < r_S$ with nothing further asked of the spillover benefit. The binding restriction of the pair is therefore (A.31), which is a lower bound on the initial premium η; (A.32) only rules out a floor so high that the convenience yield never materializes.

Remark A.2 (Location of the Welfare Trough)**.** *The trough occurs strictly before the convenience yield equalizes the two rates, $\theta^v < \eta/\nu$. At $\theta = \eta/\nu$ the EC bond rate equals the safe rate, so $\bar{q}(\eta/\nu) =$*

q_S, and with $\bar{q}'(\eta/\nu) = \nu/(1+r_S)^2$ the marginal EC cost is $\Lambda'(\eta/\nu) = r_S/(1+r_S) - \eta/(1+r_S)^2$. Evaluating the reference weight of (a′) at the knife-edge value $E^{all}[\Pi_{St}] = (1+r_S)/r_S$, which $\beta_S = q_S$ delivers when S's debt choice is interior and its marginal utility is a martingale, the marginal fee is worth $4\big[1 - \eta/(r_S(1+r_S))\big]$ against a marginal grant of 4, so

$$\left.\frac{d\mathcal{W}_S}{d\theta}\right|_{\eta/\nu} \propto \frac{4\eta}{r_S(1+r_S)} + B'(\eta/\nu) > 0, \tag{A.35}$$

and assumption (a′) alone carries the sign through from F to $F+\mathcal{R}$.[31] *The scale-economy term $-\theta\,\bar{q}'(\theta)$ in the marginal fee therefore begins recovering S's welfare while EC bonds are still more expensive than national bonds, and the spillover benefit $B'(\eta/\nu) \geq 0$ reinforces the recovery. The convenience yield accelerates it but is not necessary to initiate it.*

Multiplicity in the scale of integration. The valley shape is also what delivers the disconnected participation set of Proposition 5.2, stated in Section 5.3. Because the proof rests entirely on the single-trough geometry established here, it is collected in this subsection rather than with the baseline proofs.

Proof of Proposition 5.2. Part (i): $\Delta\mathcal{W}_S(0) = 0$ by definition. Proposition A.5(i) makes $d\mathcal{W}_S/d\theta$ negative on a right-neighborhood of 0, so $\Delta\mathcal{W}_S(\theta) < 0$ for θ small and positive. Proposition A.5(ii)–(iii) delivers a *unique* trough θ^v, with

$$\frac{d\mathcal{W}_S}{d\theta} < 0 \text{ on } (0,\theta^v), \qquad \frac{d\mathcal{W}_S}{d\theta} > 0 \text{ on } (\theta^v,\theta^\dagger] \text{ and above } \theta^\dagger,$$

so $\Delta\mathcal{W}_S$ is strictly decreasing and then strictly increasing past the trough over the whole of $[0,\theta^{\max}]$. By the maintained hypothesis $\Delta\mathcal{W}_S(\theta) > 0$ for some $\theta \leq \theta^{\max}$; continuity and the single-trough shape then imply $\Delta\mathcal{W}_S$ crosses zero exactly once from below, at a point $\theta^{\rm rec} \in (\theta^v,\theta^{\max})$, with $\Delta\mathcal{W}_S < 0$ on $(0,\theta^{\rm rec})$ and $\Delta\mathcal{W}_S \geq 0$ on $[\theta^{\rm rec},\theta^{\max}]$. Hence $\mathcal{F} = \{0\} \cup [\theta^{\rm rec},\theta^{\max}]$, which is disconnected because the excluded interval $(0,\theta^{\rm rec})$ is nonempty.

Part (ii): under the protocol a marginal expansion $d\theta > 0$ is adopted only if it leaves both countries weakly above the withdrawal payoff $\mathcal{W}_i(0,\tau)$, the tax-only union that abandoning common debt at an unchanged tax delivers. At $\theta = 0$, any small $d\theta > 0$ yields $\Delta\mathcal{W}_S < 0$ by Part (i), so no expansion is adopted and $\theta = 0$ is a rest point. On the upper component $[\theta^{\rm rec},\theta^{\max}]$ every scale satisfies $\Delta\mathcal{W}_S \geq 0$, so S participates throughout; since $\Delta\mathcal{W}_R$ is strictly increasing by hypothesis, R

[31] At the same reference weight the slack in (A.31) is $4\eta/[(1+a)r_S] - B'(0^+)$ with $a \equiv r_S + \eta$, and $F(\eta/\nu)$ exceeds it by $4\eta^2/[r_S(1+r_S)(1+r_S+\eta)] + B'(\eta/\nu) + B'(0^+) > 0$, so $|\mathcal{R}(\eta/\nu)| < F(\eta/\nu)$ by (A.30); only the uniform bound is used, not the monotonicity of $\mathcal{R}$. For an arbitrary reference weight the comparison needed is $E^{\rm all}[\Pi_{St}]\big(\Lambda'(0^+) + \Lambda'(\eta/\nu)\big) < 2 + \frac{1}{4}\big(B'(0^+) + B'(\eta/\nu)\big)$, which the knife-edge value satisfies with slack. A weight far enough above it, which is what a frequently binding barrier on S would produce, can reverse it.

prefers the largest feasible scale and the protocol comes to rest at $\theta^{\text{high}} = \theta^{\max}$, where no further expansion is possible. Both $\theta = 0$ and $\theta^{\max}$ are therefore stable.

Part (iii): The monotonicity hypothesis and Part (i) give

$$\Delta\mathcal{W}_R(\theta^{\max}) > 0, \qquad \Delta\mathcal{W}_S(\theta^{\max}) \geq 0,$$

where the first inequality also uses $\Delta\mathcal{W}_R(0) = 0$. Hence $\theta^{\max}$ weakly Pareto-dominates $\theta = 0$, strictly for R. Now consider any continuous path

$$\gamma : [0,1] \to [0, \theta^{\max}], \qquad \gamma(0) = 0, \qquad \gamma(1) = \theta^{\max}.$$

By the intermediate value theorem it attains some value in $(0, \theta^{\text{rec}})$, where $\Delta\mathcal{W}_S < 0$ and S withdraws. No participation-feasible incremental path therefore connects the two components, and reaching $\theta^{\max}$ requires a discrete jump of size at least θ^{rec}. □

A.7 Computational appendix

This subsection records how the equilibrium of Definition 2.1 is computed at each point of the policy grid. All figures and tables in the paper come from the fine specification described here.

Income. The bivariate AR(1) (2.1)–(2.2) is discretized by vector Rouwenhorst. Since $\rho_R = \rho_S$, the innovation covariance Σ is diagonalized, the Rouwenhorst method is applied to each independent component, the two transition matrices are combined by Kronecker product, and the grid is rotated back and rescaled to unit unconditional mean. Rouwenhorst rather than Tauchen is used because the process is quarterly and highly persistent, $\rho = 0.95$, a region in which Tauchen's quadrature degrades. The node counts belong to the *rotated components*, not to the two countries: 25 nodes for the higher-variance component and 15 for the other, so the joint income state space has $25 \times 15 = 375$ points. Because $\rho_R = \rho_S$ and the innovations are positively correlated, the principal axes are the sum and the difference of the two log incomes, so each of y_R and y_S is a combination of both components and takes a distinct value at every one of the 375 joint states rather than at 25 or 15 of them. The first two unconditional moments and the autocorrelations of the discretized chain are checked against their continuous counterparts.

Debt grids. Debt is discretized on evenly spaced grids. The risky country's grid runs from the savings allowance to the borrowing limit, $b_R \in [-\underline{a}, \bar{b}_R] = [-0.75, 7.0]$ in units of one quarter's output, on 311 nodes. The safe country's *choice* set is the band $[\underline{b}_S, \bar{b}_S] = [0.80, 4.16]$ on 97 nodes. Its *state* grid is wider, extending a full date-0 grant beyond the band at that same spacing, to -1.230 on 155 nodes, so that the translated argument $b_S - 4\theta$ in (3.2) always lies inside the grid at

every θ on the policy grid and no occupiable state has to be extrapolated. Because no admissible policy can leave the band, the extra nodes carry zero ergodic mass by construction. The risky country's grid is deliberately not extended in the same way, because R's ergodic distribution sits deep in borrowing territory and widening its savings region would raise the cost of its value-function iteration for states it never visits. A small tail of R's mass therefore does translate past the top of its grid at large θ: none at $\theta \leq 0.4$ and a fraction of a percent of the in-market mass at the top of the policy grid. Those query points are clamped to the top node and credited with the utility of consuming the residual they cannot carry as assets, which is feasible at the translated state but not optimal there and so is a lower bound. Measured against a solution on a grid wide enough to avoid the clamp entirely, the treatment costs less than one percent of the welfare differences being reported.

Solution. The equilibrium is block recursive, so R is solved first. Its problem is a value-function iteration with Howard policy improvement, nested inside an outer loop that resets the price schedule to the break-even price implied by the current default rule. The price update is undamped, so the sup-norm residual on q_R is the residual of the price map itself and convergence is a genuine fixed-point certificate; the operator is monotone but not a contraction, so the solver falls back permanently to a damped update if the residual fails to improve for 100 consecutive iterations. Because $\gamma_i > 1$ makes flow utility negative, the cold start $V = 0$, $q_R = 1/(1 + r_R)$ is an exact upper bound and the iteration descends from above. That guarantee lapses under the public-good extension of Section 5.1, where the additive term $\vartheta \log G_t$ can be positive, so the price operator there has no descent property; convergence on that grid is correspondingly less clean. Three of the 441 public-good points fall back to the damped update and two stop at the iteration cap with value residuals of order 10^{-3} and price residuals of order 10^{-2}, against 10^{-6} everywhere on the baseline grid; the implied error is a few hundredths of a consumption-equivalent percentage point, which is immaterial at the magnitudes reported in Section 5.1 but is the reason the public-good surface is visibly rougher than the baseline one. The value of a defaulted claim, Φ in (2.20), is obtained in a single downward sweep of the debt grid: since $|(1 - h)b| < |b|$, the value at any node depends only on nodes nearer zero. Country S is then solved by value-function iteration taking R's default and re-entry probabilities as given, with two value functions according to whether R has market access. Convergence is to a sup-norm tolerance of 10^{-6} with a cap of 1200 outer iterations; Howard steps per outer iteration are 20 for R and 200 for S, the latter larger because β_S is close to one and the effective decay rate of S's iteration is correspondingly slower.

Distributions and moments. Ergodic distributions are computed analytically, as the invariant measure of the Markov chain that the equilibrium policies induce on the discretized state space, rather than by simulation. For R the chain has two pools, in-market and post-default exclusion,

with flows between them governed by the default rule and the re-entry probability α; for S it has two regimes according to R's market access. Welfare integrals in (3.2) are taken against these measures, with the shifted argument $b_i - 4\theta$ evaluated by shape-preserving cubic interpolation along the debt grid. Reported moments are placed on an annual basis: default frequencies and spreads are scaled to annual rates, and debt ratios are divided by four so that the denominator is annual rather than quarterly output. The default frequency is unconditional, the share of the entire ergodic mass, in-market and excluded pools together, that defaults per period; conditioning instead on market access, as is also common, would raise it by the excluded share of about six percent of its own value. The debt ratio and the mean spread are averages over the in-market pool and over repaying in-market states respectively, since debt outstanding and the yield on newly issued debt are defined only there.

Channel derivatives. The three equilibrium-response channels of Section 4.3 are evaluated by least squares over a symmetric window of policy points spanning ten percent of each axis, rather than by adjacent differences. These are the two price responses $\partial q_R/\partial\theta$ and $\partial q_R/\partial\tau$ and the induced change in R's default frequency. The price schedule and R's default probabilities are piecewise constant in the policy, so a difference against a single neighbor returns either exactly zero or a jump divided by the grid step, whereas a window spreads each discrete flip over its width. The estimates are correspondingly imprecise in the immediate neighborhood of $\tau = 0$, where R's default set moves discretely.

Policy grid. The surfaces in Section 4 are computed on a 21×21 grid, evenly spaced over $\theta \in [0,\, 0.5]$ and $\tau \in [0,\, 0.02]$, so each of the 441 cells is a separate recursive equilibrium. The step in θ is 0.025, which places the reference program $\theta = 0.1$ exactly on a node. The public-good extension of Section 5.1 repeats the exercise on a grid of the same size with $\tau \in [0,\, 0.05]$.

A.8 Additional figures

Figure A.3: Ergodic distribution of the debt-to-output ratio at the no-integration benchmark

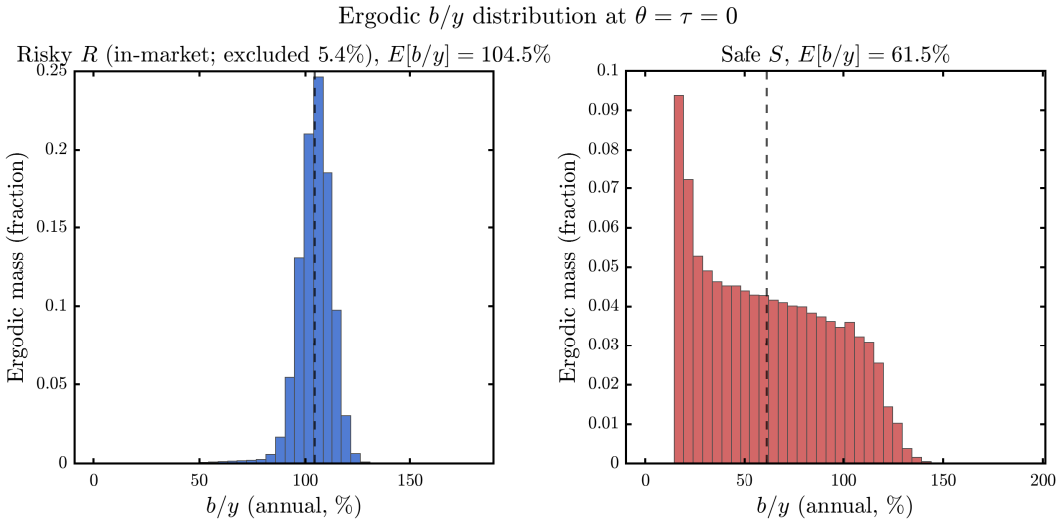

The figure shows the ergodic debt distributions that lie behind the two debt moments the calibration targets. Each panel is a mass-weighted histogram of the annual debt-to-output ratio b/y in the ergodic distribution at the no-integration benchmark $(\theta, \tau) = (0, 0)$, debt being plotted as a positive liability stock. *Left*: risky country R, restricted to the in-market pool over which $E[b_R/y_R]$ is averaged; the share of ergodic mass in post-default exclusion is noted in the panel title. That pool is excluded from the histogram rather than placed at $b = 0$: under partial recovery an excluded sovereign still carries its defaulted stock and re-enters owing $(1-h)b_R$, so its debt is not zero and is not priced in the market. *Right*: safe country S. The dashed vertical line in each panel marks the reported mean $E[b/y]$. The broad spread of S's distribution reflects the deliberate $\beta_S = q_S$ knife-edge calibration and the reflecting band of Section 2: with marginal utility a martingale, S's debt follows a driftless random walk between reflecting barriers, so its mean is pinned near the midpoint of the pair $(\underline{b}_S, \bar{b}_S)$ while its dispersion is left unrestricted. The distributions are the invariant measures of the Markov chains induced by the equilibrium policies, computed analytically rather than by simulation, at the parameter values of Table 4.1.

Figure A.4: Ergodic distribution of the sovereign spread at the no-integration benchmark

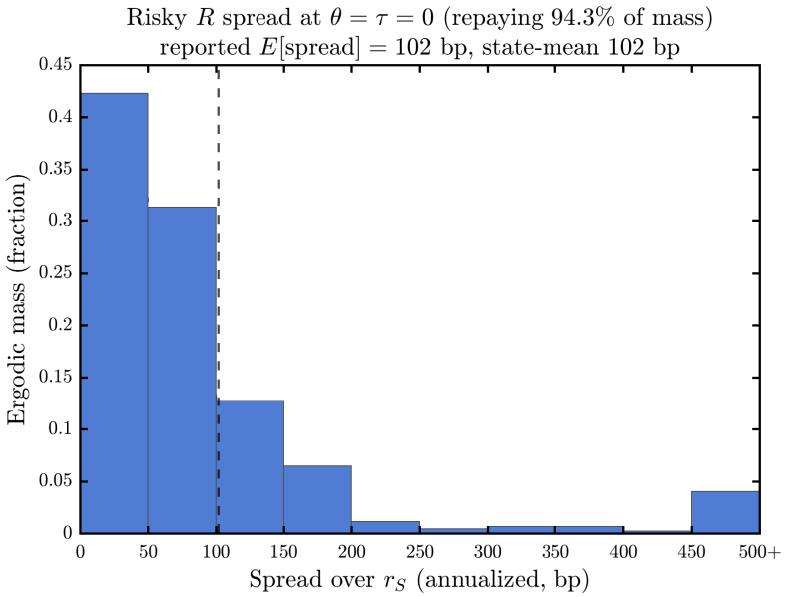

The figure shows the ergodic distribution behind the mean spread the calibration targets, and how far that mean sits from the state the risky country typically occupies. It is a mass-weighted histogram of the risky country's annualized sovereign spread over the safe rate r_S at the no-integration benchmark $(\theta, \tau) = (0, 0)$. Only R is shown: the safe country never defaults and borrows at the constant risk-free price q_S, so its spread is identically zero. The distribution is restricted to states in which R repays and retains market access, the same set over which the reported mean spread is computed, since the spread is the yield on debt actually issued and is undefined in default and exclusion; the repaying share of ergodic mass is noted in the panel title. Each state's spread is annualized geometrically as $(1/q_R)^4 - (1 + r_S)^4$ and reported in basis points, in bins of 50 bp; the final bin is top-coded, so mass beyond 500 bp is collected at the right edge and the visible tail understates the most extreme states. The dashed vertical line marks the reported mean spread, which annualizes the ergodic mean of the quarterly spread and is given in the panel title alongside the mean of the annualized state-by-state spreads plotted here, the two differing by a Jensen gap of a few basis points. The distribution is markedly right-skewed: the modal state is priced close to risk-free, and the reported mean is pulled above it by a thin tail of high-debt, low-income states in which default risk is imminent. The mean spread is therefore not representative of the spread R faces in typical times, and the moment is informative about the tail of the ergodic distribution rather than about its center. The distribution is the invariant measure of the Markov chain induced by the equilibrium policies, computed analytically rather than by simulation, at the parameter values of Table 4.1.

Figure A.5: Equilibrium bond-price schedule of the risky country

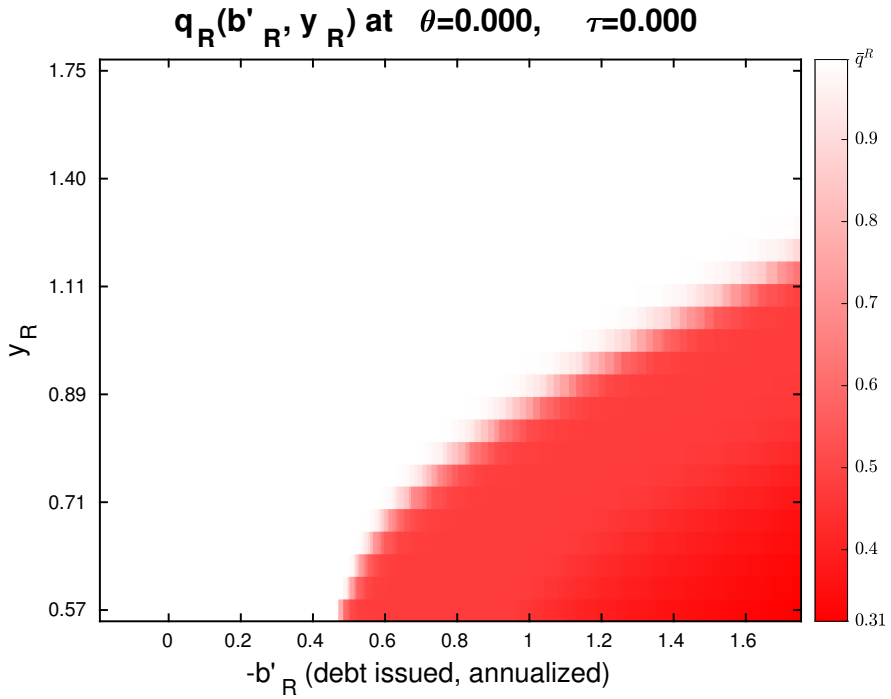

The figure shows the price schedule that private lenders offer the risky country, the equilibrium object from which both its spread and its default frequency follow. It plots the private-market bond price $q_R(b'_R, y_R)$ at the no-integration benchmark $(\theta, \tau) = (0, 0)$, as a function of next-period debt b'_R, expressed as a positive liability on an annual-output basis (horizontal axis), and the risky country's income y_R (vertical axis). The plotted coordinate $-b'_R$ reverses MATLAB's $b'_R \leq 0$ borrowing convention; elsewhere in the paper, $b'_R \geq 0$ denotes the same liability. Because y_R takes a distinct value at each of the 375 joint income states, the vertical axis is not a model grid: the joint states are binned into 25 equal intervals of $\log y_R$, and within each bin the price is the ergodic-probability-weighted average of q_R over the states it contains. Darker red indicates a lower price (a higher implied spread and default risk): the price approaches R's own risk-free price $\bar{q}^R \equiv 1/(1 + r_R)$ at low debt, which is distinct from the EC issue price $\bar{q} = 1/(1 + r_S)$, and falls to about a third of par at the largest issuance the grid allows. The color scale runs from the lowest price actually attained rather than from zero, so it spends its range on the prices R faces. Computed from the solved recursive equilibrium at the parameter values of Table 4.1.

Figure A.6: Robustness of S's welfare to the spillover cost φ

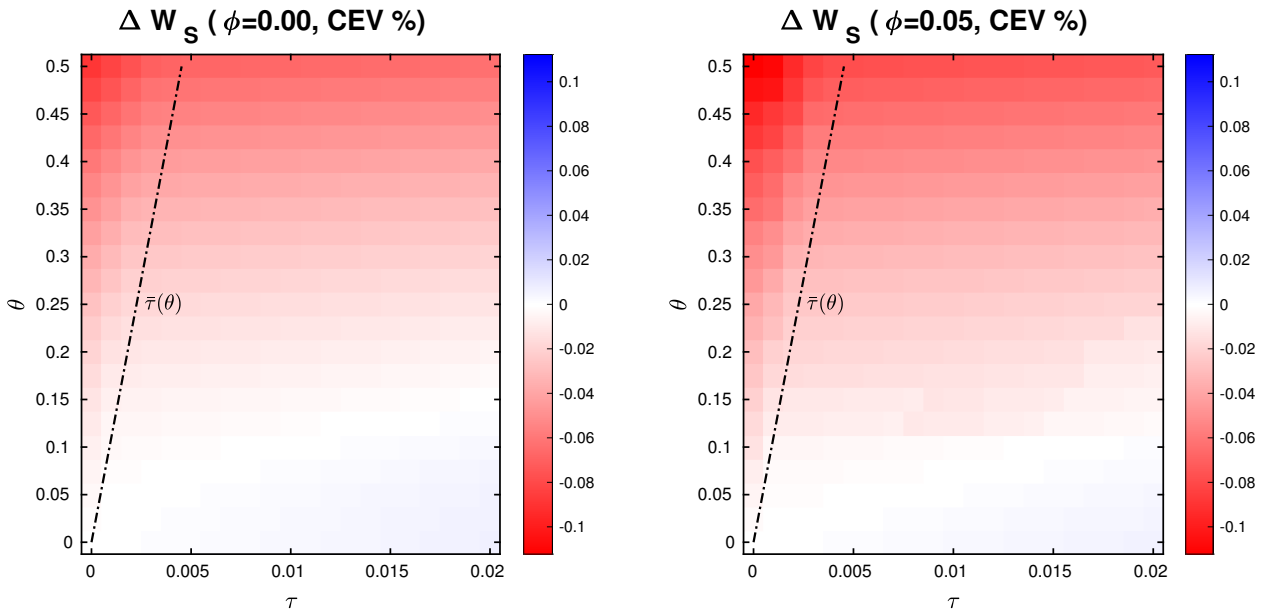

The figure shows how sensitive the safe country's welfare surface is to the one parameter with no direct empirical counterpart, the spillover cost φ it bears while the risky country is excluded from markets. Each panel plots the consumption-equivalent welfare change of S, in percent, over the same 21×21 policy grid as Figure 4.1, holding every other parameter at the baseline of Table 4.1. *Left*: $\varphi = 0$. *Right*: $\varphi = 0.05$. The baseline value is $\varphi = 0.02$, so the two panels bracket it. Because φ enters only S's problem and is absent from R's value function, country R's welfare surface is unchanged across these panels and is therefore not reproduced. Blue indicates a welfare gain, red a loss. The dash-dotted line marks the fiscal threshold $\bar{\tau}(\theta)$, which does not depend on φ. Both panels share a common color scale, so shading is comparable across them. Each panel is drawn from 441 separately solved recursive equilibria.

Figure A.7: Changes in average debt-to-output ratios relative to no integration

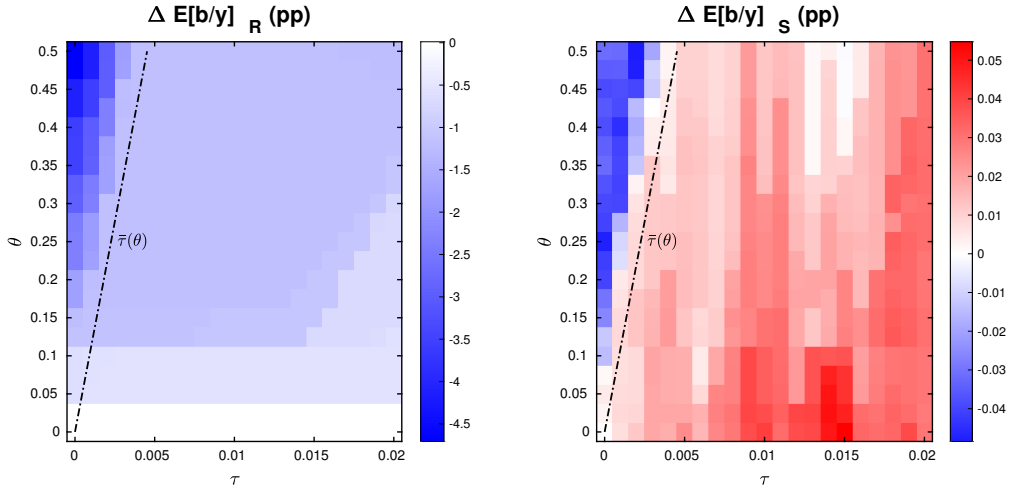

The figure shows how much national debt each country carries under the common program, relative to the no-integration point $(\theta, \tau) = (0, 0)$, at every cell of the 21×21 policy grid. *Left*: change in R's mean ergodic debt-to-output ratio $E[b_R/y_R]$, in percentage points. *Right*: the same for S. Both ratios are on an annual basis, the denominator being four quarters of output. Blue indicates a reduction in debt, red an increase, on color scales set separately for the two panels. The dash-dotted line marks the fiscal threshold $\bar{\tau}(\theta)$. Computed from the same solved equilibria as Figure 4.1, at the parameter values of Table 4.1.

Figure A.8: Changes in default probability and sovereign spread

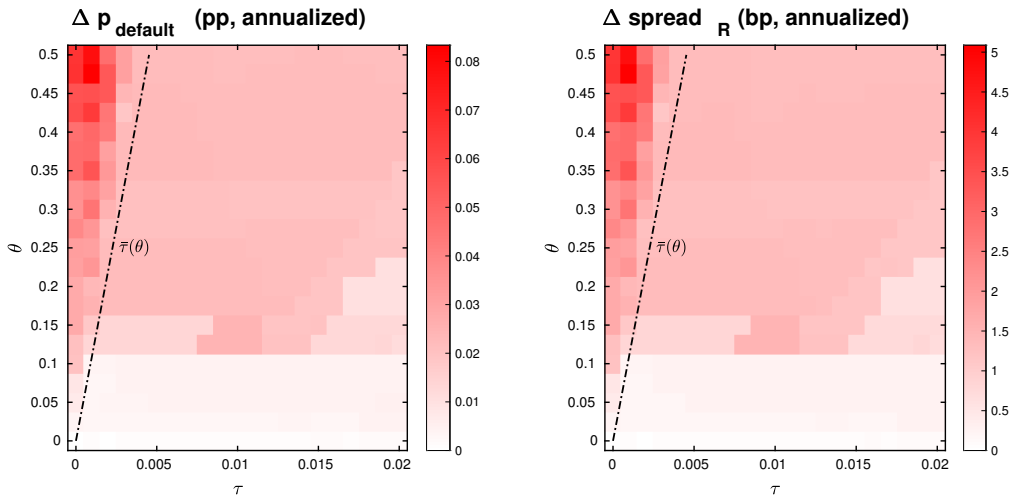

The figure shows what the common program does to the risky country's credit, at every cell of the 21×21 policy grid. Both panels are on an annual basis. *Left*: change in R's default frequency, in percentage points per year, obtained by scaling the per-quarter default probability by four. *Right*: change in R's mean sovereign spread, in annualized basis points. At each policy cell the ergodic mean of the quarterly spread is annualized geometrically, by compounding R's mean gross quarterly yield and the safe gross rate over four quarters and differencing; this is the annualization of a cell mean, not the cell mean of the state-by-state annualizations of Figure A.4. Changes are measured against the no-integration point $(\theta, \tau) = (0, 0)$. Blue indicates a reduction, red an increase, on color scales set separately for the two panels. The dash-dotted line marks the fiscal threshold $\bar{\tau}(\theta)$. Computed from the same solved equilibria as Figure 4.1, at the parameter values of Table 4.1.

Figure A.9: Welfare decomposition with respect to θ: $d\mathcal{W}_R/d\theta$ (left) and $d\mathcal{W}_S/d\theta$ (right)

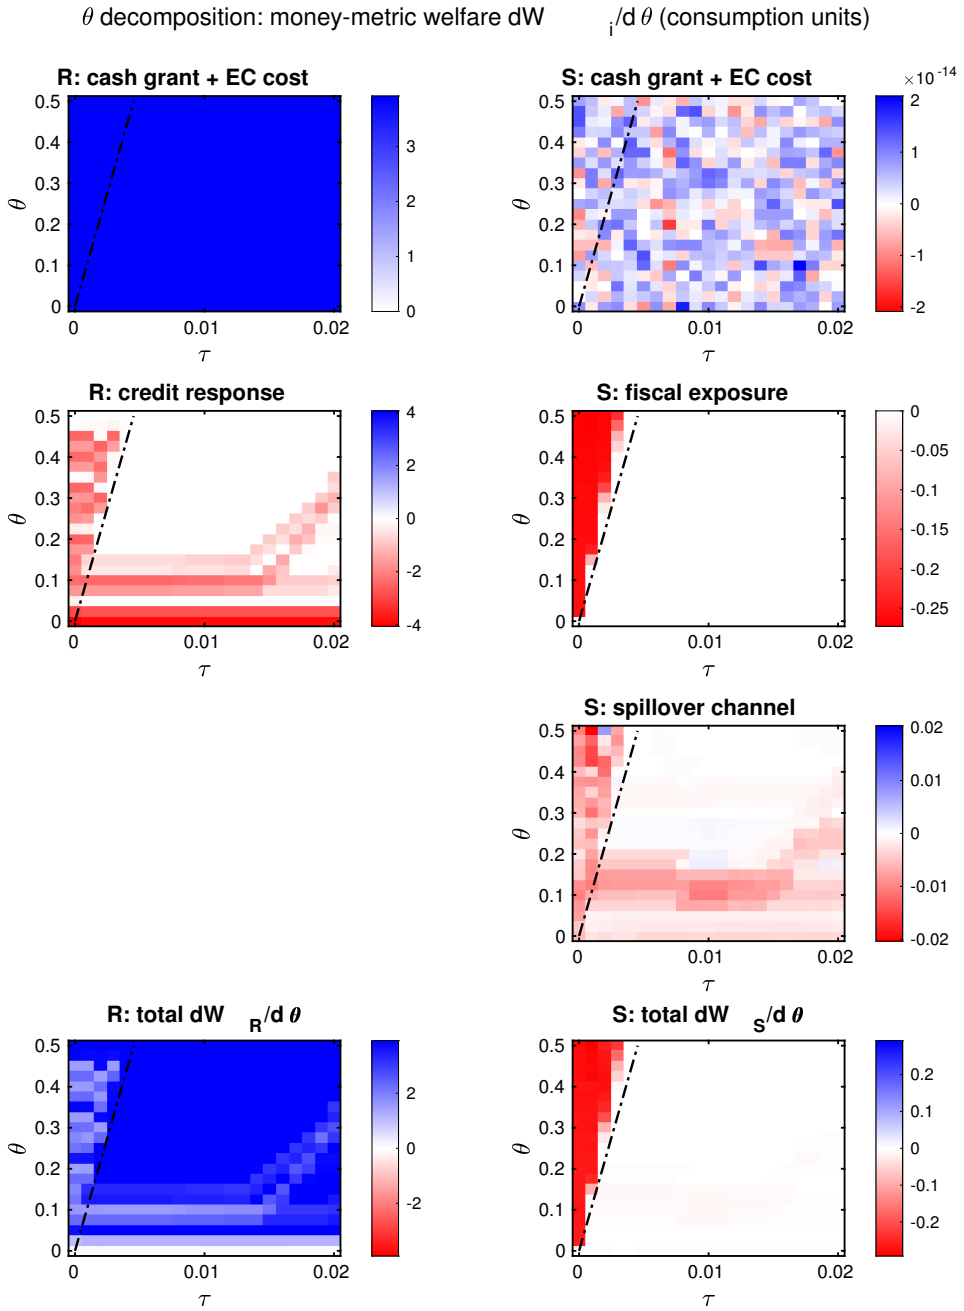

The figure shows every channel of the θ decompositions over the whole policy grid, of which Table 4.2 reports one cell and Figure 4.2 one horizontal slice. The left column shows the channels of R's decomposition (3.8) and the right column those of S's (3.10), on the same 21×21 grid as Figure 4.1, with the total $d\mathcal{W}_i/d\theta$ in the bottom row. R has one channel fewer than S, so the third panel of the left column is blank. Units are date-0 consumption per unit of the instrument, the scale of (3.5) rather than the consumption-equivalent percentages of Figure 4.1. Blue indicates a positive contribution to welfare, red a negative one, on a color scale set panel by panel, since the channels differ by orders of magnitude. The dash-dotted line marks the fiscal threshold $\bar{\tau}(\theta)$; the threshold-contingent channel, S's fiscal exposure, vanishes to its right. Computed from the same solved equilibria as Figure 4.1, at the parameter values of Table 4.1.

Figure A.10: Welfare decomposition with respect to τ: $d\mathcal{W}_R/d\tau$ (left) and $d\mathcal{W}_S/d\tau$ (right)

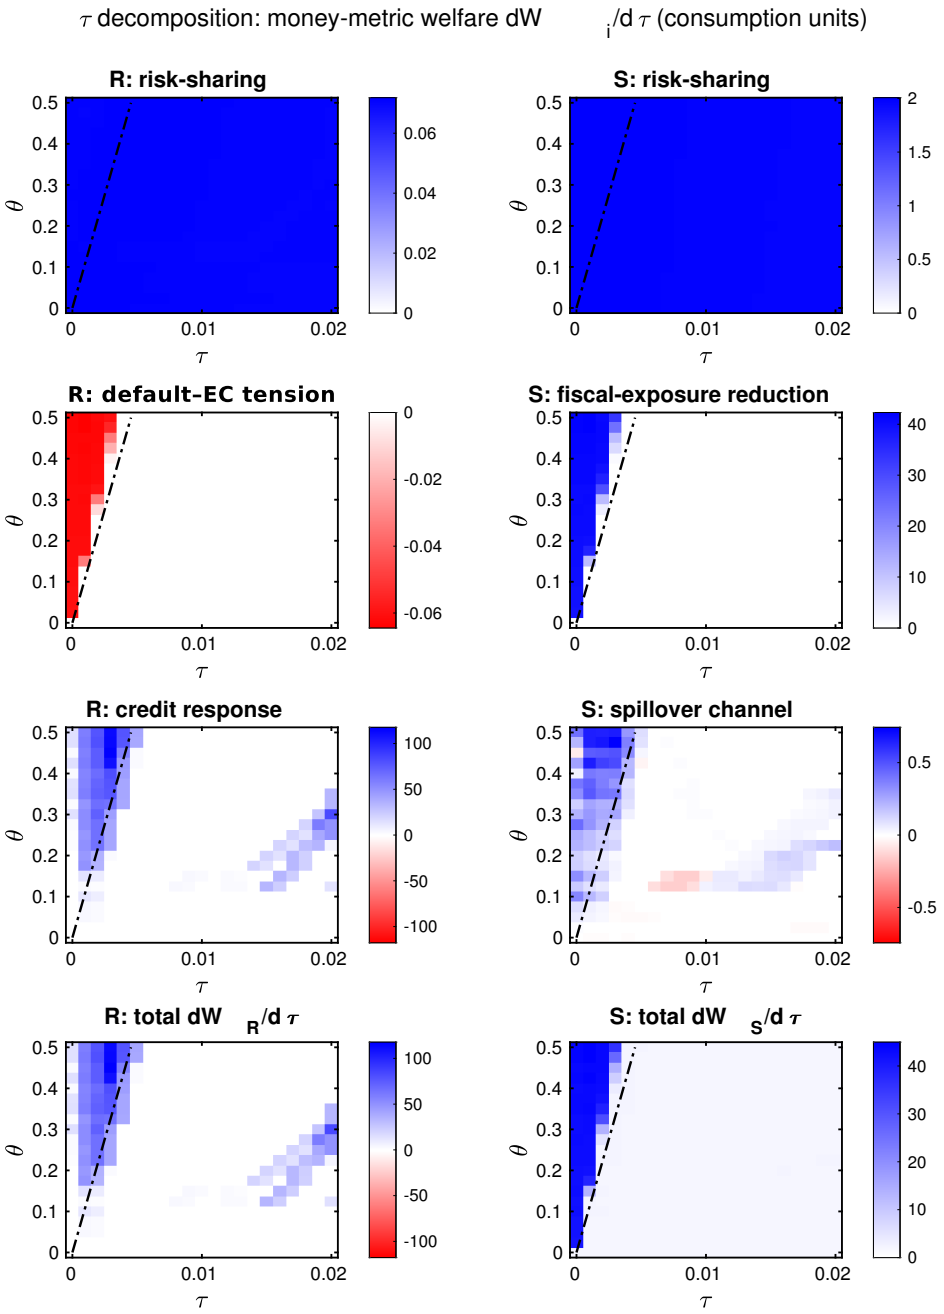

The figure shows every channel of the τ decompositions over the whole policy grid, the counterpart of Figure A.9 for the tax instrument. The left column shows the channels of R's decomposition (3.13) and the right column those of S's (3.15), on the same 21×21 grid as Figure 4.1, with the total $d\mathcal{W}_i/d\tau$ in the bottom row. Units are date-0 consumption per unit of the instrument, the scale of (3.5) rather than the consumption-equivalent percentages of Figure 4.1. Blue indicates a positive contribution to welfare, red a negative one, on a color scale set panel by panel, since the channels differ by orders of magnitude. The dash-dotted line marks the fiscal threshold $\bar{\tau}(\theta)$; the threshold-contingent channels, default–EC tension for R and fiscal-exposure reduction for S, vanish to its right. Computed from the same solved equilibria as Figure 4.1, at the parameter values of Table 4.1.

Figure A.11: Welfare changes evaluated at each policy's own invariant distribution

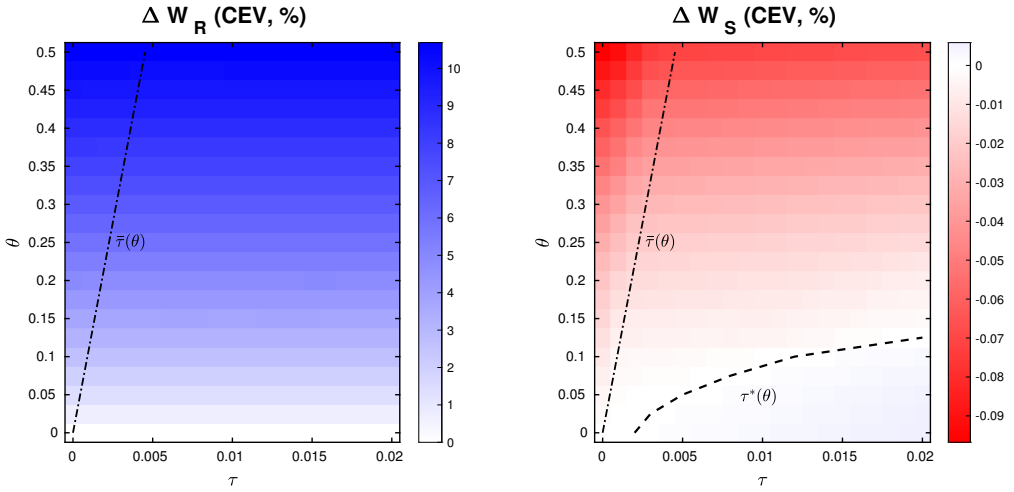

The figure shows how much of the headline welfare result depends on the distribution from which the reform is launched. It plots the consumption-equivalent welfare change, in percent, under $\mathcal{W}_i^{\text{own}}$ of (A.12): the same date-0 reform as the headline measure, grant included, but launched from each policy's own invariant distribution $\mu^*(\cdot;\theta,\tau)$ instead of the pre-reform one. The left panel is the risky country R and the right panel the safe country S; panels, curves, color convention and the 21×21 policy grid are otherwise as in Figure 4.1, with which this figure should be compared cell by cell. The two differ only in the initial distribution, so the gap between them is the displacement (A.13) and nothing else. This is not a stationary comparison, which is Figure A.12 instead. Computed from the same solved equilibria as Figure 4.1, at the parameter values of Table 4.1.

Figure A.12: Welfare changes in the mature regime (stationary comparison)

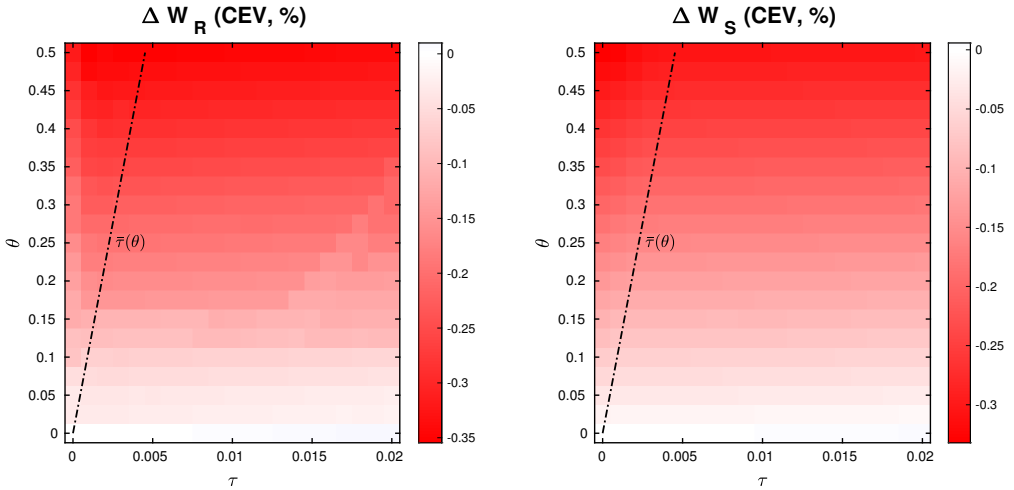

The figure shows what the arrangement is worth to a member that has lived under it long enough for the one-time transfer to have washed out of the distribution. It plots the consumption-equivalent welfare change, in percent, under $\mathcal{W}_i^{\text{ss}}$ of (A.14), the conventional stationary comparison, computed on the full 21×21 policy grid: the value function is integrated against each policy's own invariant distribution over its full support, with no date-0 grant and no conditioning on R having market access. The left panel is the risky country R and the right panel the safe country S; panels and policy grid are as in Figure 4.1, and the dash-dotted curve is again the fiscal threshold $\bar{\tau}(\theta)$; note the color scale differs by more than an order of magnitude from that figure. The compensating tax $\tau^*(\theta)$ is not drawn, because it is the tax that compensates S for adopting a program and so belongs to the launch experiment, which this measure is not. Because an invariant distribution has forgotten a one-time transfer, this measure prices the perpetual EC fee and the ongoing tax while excluding the resources they fund, so it is not an assessment of adopting the program and is not comparable cell by cell with Figure 4.1. It answers what a member faces once the transfer is history. Computed from the same solved equilibria as Figure 4.1, at the parameter values of Table 4.1.

Figure A.13: Public-good provision in the mature regime (stationary comparison)

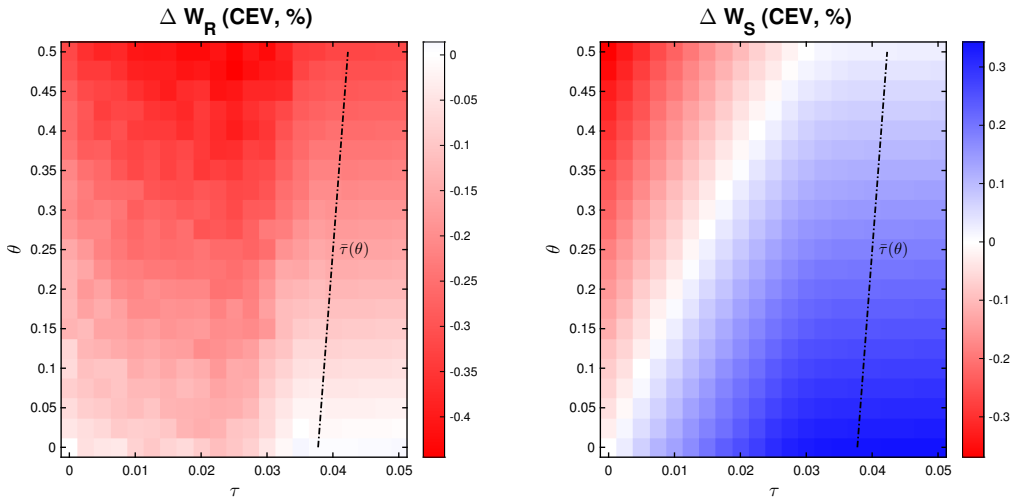

The figure shows what the public-good arrangement of Section 5.1 is worth to a member that has lived under it long enough for the one-time transfer to have washed out of the distribution. It plots the consumption-equivalent welfare change, in percent, under $\mathcal{W}_i^{\text{ss}}$ of (A.14), for the economy of Figure 5.1 and on that economy's 21×21 grid with $\tau \in [0, 0.05]$: no date-0 grant, the full invariant support, and no conditioning on R having market access. The left panel is the risky country R and the right panel the safe country S, on independent color scales. Both are measured against $(\theta, \tau) = (0, 0)$ *within* the public-good model, so the reference point already carries the expenditure ξ. The dash-dotted curve is the public-good threshold $\bar{\tau}^{\,\xi}(\theta)$; the compensating tax $\tau^*(\theta)$ is not drawn, for the reason given under Figure A.12. This is the public-good counterpart of Figure A.12 and inherits its caveats: it prices the perpetual fee and the ongoing tax while excluding the one-time resources they fund, so it is not an assessment of adopting the program and is not comparable cell by cell with Figure 5.1. The safe country turns positive at a tax between 7 and 77 percent of $\bar{\tau}^{\,\xi}(\theta)$, depending on θ, so the threshold reversal of Section 5.1 survives into this measure. Along the $\theta = 0$ row $\Delta\mathcal{W}_R^{\text{ss}}$ is not monotone in τ, wandering by up to 0.05 percentage points, so the seven Pareto-positive cells in that row sit inside the resolution of the surface and the sign of R's gain there should not be read as established. Computed from the same solved equilibria as Figure 5.1, at the parameter values of Table 4.1 with $\xi = 0.05$.

Acknowledgements

I gratefully acknowledge comments from Eduardo Dávila, Tano Santos, Stephanie Schmitt-Grohé, Jesse Schreger, Martín Uribe, and Stijn van Nieuwerburgh. All remaining errors are my own.
This paper was prepared under the Lamfalussy Fellowship Programme sponsored by the ECB.

Giovanni Bonfanti
National Bureau of Economic Research (NBER), Cambridge, Massachusetts, United States; email: giovanni.bonfanti@columbia.edu

© European Central Bank, 2026

Postal address 60640 Frankfurt am Main, Germany
Telephone +49 69 1344 0
Website www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from www.ecb.europa.eu, from the Social Science Research Network electronic library or from RePEc: Research Papers in Economics. Information on all of the papers published in the ECB Working Paper Series can be found on the ECB's website.

PDF ISBN 978-92-899-8029-6 ISSN 1725-2806 doi:10.2866/5107881 QB-01-26-235-EN-N